Enhanced Multi-Modal Access to Planetary Exploration

Norm Lamarra, Richard Doyle, Jay Wyatt @jpl.nasa.gov
Jet Propulsion Laboratory/Caltech
Pasadena, CA 91109
Norman.Lamarra@jpl.nasa.gov 818-393-1561

Abstract— Tomorrow's Interplanetary Network (IPN) will evolve from JPL's Deep-Space Network (DSN) and provide key capabilities to future investigators, such as simplified acquisition of higher-quality science at remote sites and enriched access to these sites. These capabilities could also be used to foster public interest, e.g., by making it possible for students to explore these environments personally, eventually perhaps interacting with a virtual world whose models could be populated by data obtained continuously from the IPN. Our paper looks at JPL's approach to making this evolution happen, starting from improved communications. Evolving space protocols (e.g., today's CCSDS proximity and file-transfer protocols) will provide the underpinning of such communications in the next decades, just as today's rich web was enabled by progress in Internet Protocols starting from the early 1970's (ARPAnet research). A key architectural thrust of this effort is to deploy persistent infrastructure incrementally, using a layered service model, where later higher-layer capabilities (such as adaptive science planning) are enabled by earlier lower-layer services (such as automated routing of object-based messages). In practice, there is also a mind shift needed from an engineering culture raised on point-to-point single-function communications (command uplink, telemetry downlink), to one in which assets are only indirectly accessed, via well-defined interfaces. For example, instead of sending a specific instrument command sequence, a higher-level science goal could be requested; this would be automatically broken down into prioritized requests for progressively lower-level resources, some of which may not be on-board. In the downlink direction, instead of sending only raw data (which is limited by direct bandwidth), higher-value information could be gathered, perhaps using several sensors, and delivered automatically to a workstation under the scientist's control. We are thus aiming to foster a "community of access" both among space assets and the humans who control them. This enables appropriate (perhaps eventually optimized) sharing of services and resources to the greater benefit of all participants. We envision such usage to be as automated in the future as using a cell phone is today – all the steps in creating the real-time link are automated (including identification, location, routing, and of course billing!).

TABLE OF CONTENTS

1. INTRODUCTION

JPL envisions evolution of the Deep Space Network (DSN) into an Interplanetary Network (IPN) [1,2] that will provide electronic accessibility to deep space in much the same way that almost all of Earth is accessible today via the terrestrial Internet. There is untapped potential for a new level of space applications that could provide untold benefits, just as we are still discovering the vast potential for wired and wireless Internet applications. For example, the CEO of the world's dominant PC software company has described the benefits of a "connected home" (including the one he inhabits) in which individual devices, from lights to appliances, can be queried and commanded, not just locally, but remotely from our offices or automobiles. JPL is developing the implications of extending a vision like this to exploration across our solar system, for example considering human artifacts in deep space as nodes on the IPN. A primary goal is thus to make it as easy for future scientists to interact with extremely remote environments, as for future consumers to interact with their networked world. Spacecraft such as orbiters, landers, rovers and other robots can thus be viewed as potentially cooperating distributed resources, albeit ones with certain differences from those on or near Earth. Over time, a permanent community of such connected robots could be established, initially focused on enriching the exploration of other planets, but evolving to assist in making human habitation possible for some of them. The IPN is therefore a critical part of moving that vision from Science Fiction into Technological Fact.

A complementary conceptual evolution has also been espoused for some years by others at JPL, in which instruments participate in a "sensor web" that can exhibit aggregate behavior that is surprisingly rich considering the simplicity of the individual nodes [4]. Indeed, there are already terrestrial deployments of such webs, including remote ones in the Arctic and on the ocean floor, performing automated monitoring of physical processes that can affect our world dramatically, such as seismic and geothermal events. Such webs provide successful examples of remote exploration in places that are too difficult, dangerous, or costly for humans; moreover, they demonstrate the feasibility and utility of an automated and continuous presence, watching and waiting for important events. This can vastly extend the reach and coverage of scientists, who can gather the information without leaving their workplaces, analyze it collaboratively, draw appropriate conclusions from such events, and publish their findings and hypotheses, all over the Internet.

2. GOALS/BENEFITS

A particularly important motivation for the IPN emerges from examination of the history of exploration: a watershed event in the opening up of any frontier is when infrastructure begins to be deployed in the remote environment. Once this step is taken, no longer must expeditions or missions carry all required resources – in self-contained fashion – into the remote environment. Instead, each mission can rely on services and infrastructure already in place, and focus more completely on how to accomplish its given objectives. One example is the establishment of permanent base camps in icy regions. Another is the network of GPS satellites, originally deployed for military reasons, yet now used by autos, boats and even bicycles. A third example is the construction of highways or railways, enabling settlement following on the heels of early pioneers. We suggest that similar leaps in capability are waiting to be made in space exploration, by leveraging thoughtful investment in IPN infrastructure.

A primary underlying goal of the IPN is thus to leverage such infrastructure investment into significant improvement in science quality and quantity. One dimension of this relates to sensor capability: currently, we are often forced to drastically limit our utilization of sensors at remote planets. For example, the Mars Global Surveyor is able to return less than 1% of its possible sensor data to Earth, because of limitations such as downlink bandwidth. Compounding this problem, we expect continued rapid evolution of terrestrial sensor capability, such as image resolution, detector sensitivity, and internal processing. Future deep-space missions thus cannot fully benefit from such improvements because of current limitations in capabilities of the DSN. Ubiquitous high-speed connectivity provided by the IPN could thus significantly enhance science exploration by extending the applicability and coverage of such evolving sensors into deep space.

Another dimension of improving science is via increased sensor modality. We currently make little effort to project senses like sound, touch and smell into deep space, but all of these can be achieved electronically with appropriate interfaces. For example, physical characteristics like shape, mass and even texture can be conveyed electronically via haptic devices. It would thus be possible to construct a virtual world containing accurate models of such characteristics. Even though direct remote interactivity with a remote world is problematic, due to round-trip light time, it would nevertheless be possible for a scientist to interact dynamically with such a virtual world locally, and evolve the realism of the models by incorporating actual sensor measurements.

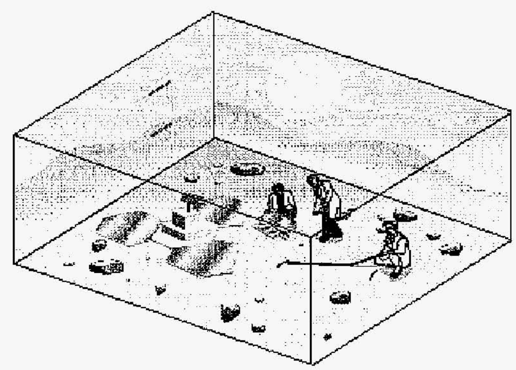

Figure 1: IPN Science Scenario – *Virtual Sense of Presence*

When coupled with the other advantage offered by the IPN – greatly enhanced bandwidth – the possibilities for how PIs interact with sensors and instruments expand dramatically, particularly if virtual environments technology comes into play. In a properly conceived immersive environment, a PI could enter a depiction of say, a Mars surface environment, and – depending on what kinds of data have been collected, returned and rendered – walk around objects, pick them up, feel their texture, etc. The PI's choices about what next observation to make, or next sample to collect, would be much more richly informed under this vision of how to exploit the power of the IPN than under today's DSN-provided bit streams coupled with comparatively low-end approaches to visualization. See Figure 1. This concept extends naturally to NASA's goal of improving public access: compare the historical impact of millions watching national TV coverage of the first moon walk to that of a few enthusiasts able to look at the surface of the moon through a backyard telescope.

In addition to the cost and efficiency advantage of using the same infrastructure for each new mission, having local communications and navigation services available in the remote environment opens up other advantages in the long term. For example, mission concepts based on having several heterogeneous assets active concurrently in the same environment – assets which can coordinate their activities via the local infrastructure, without having all such mediation make the round-trip to Earth. One straightforward scenario could be at Mars: an orbital platform detects evidence of recent water at the surface. A general request for further *in situ* investigation goes out to a fleet of mobile surface assets. After considering factors such as proximity and payload, one particular rover is selected to investigate further. See Figure 2.

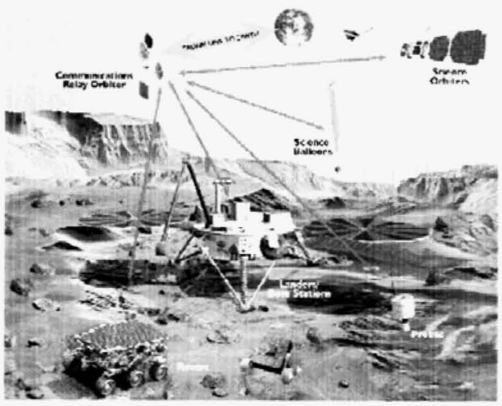

Figure 2: Coordinating Mars assets

Long before the IPN enables autonomous deep space assets to coordinate their activities, it will pay dividends by providing transparent, higher-level mission services to NASA's Principal Investigator community. Under the DSN, PIs are included, painfully, in the fine points of commanding and transporting data from spacecraft. Under the truly networked communications of the IPN, standard middleware services such as *store and forward* would hide such details from the scientist, providing for a direct link between investigator and payload – "bringing the sensors to the scientists." The final, logically straightforward step of linking the IPN with the terrestrial internet (with appropriate layers of IT security) allows the PI – and her colleagues – to enjoy this intimacy with remote assets from the convenience of the home institution.

A final example of a stretch goal for the IPN is to improve adaptability. Part of the motivation for this is that we often know little about what to look for until we get there, and then realize we didn't bring the right tools. Even well-planned deep-space science missions often undergo significant evolution in approach and goals when the data come back and interpretation begins. Again, we see a natural role for infrastructure in enabling adaptability; for example, it is much easier to modify software built from modular components plugged into an established and robust framework, than to add capability by rebuilding an entire monolithic application.

3. Approach/Technology

Many of the technologies required to achieve these goals are already in common use today on Earth. However, we need to focus on the special problems facing application of those technologies for long-term use in deep space. Such problems include: survivability in harsh environments of extreme temperature, radiation and chemical challenges; tight constraints on power, mass and volume; and of course the inescapable reality of limitations due to cost. Indeed, this last factor may potentially be one of the largest selling points of infrastructure in general and the IPN in particular. It seems obvious that the richer the electronic infrastructure, the easier it would be to achieve higher performance in many areas, and hence the more cost-effective the solution, provided the infrastructure costs do not dominate. For example, a simple sensor web deployed on Mars could achieve much larger coverage at lower cost, provided there is capability to gather, process and deliver the information provided by the sensors. Similarly, a capable rover could return much richer science if it could rely on high-speed communications to Earth, e.g., via permanent orbiting relay.

One approach toward leveraging those technologies for use in space is via space middleware services [5]. In this context, middleware is viewed as software involved in connecting separate application pieces (components), but excluding that involved in the communications link itself. Indeed, a primary goal of such middleware is to provide distributed capability that hides the vagaries of communication. Mddleware can thus be viewed as residing between application components (layered "above") and communication components (layered "below"), as shown in Figure 3.

This IPN-based approach to accomplishing science missions could be further enhanced by application-level capabilities for *in situ* science data analysis and automated mission planning – built on top of middleware services. Each PI could have "proxy" software onboard remote assets, created and evolved under their direction. Whenever sought-after or unexpected intermediate analysis results appear, or alternatively, whenever a PI gains a new insight in the immersive environment, she can issue new requests to the remote assets, with automated mission planning accommodating those requests. A science team could co-exist this way, all accessing – via the IPN – the set of deep space assets, and the remote environments they have been delivered to, as their collective laboratory.

While infrastructure such as space middleware can simplify spacecraft operation and cooperation, it becomes more obviously cost-effective as more craft are involved (and as more resources can be leveraged). For example, at this date there are two missions recently launched for

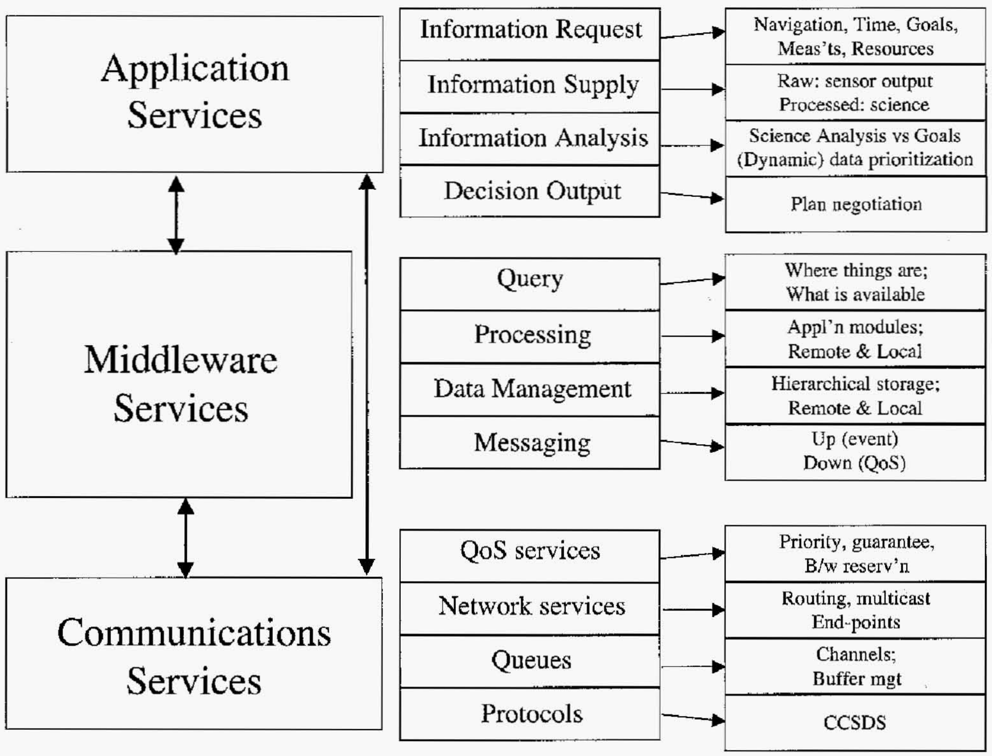

Figure 3: Layered View of Space Middleware

Mars, and three others in development, while others have already been operating in Mars orbit for some time. There are future plans to fly cooperative spacecraft, even some in tight formation. This is therefore a good time to consider the potential benefits of IPN infrastructure, initially driven by the Mars context.

Specifically, we have looked at two key enabling technologies that support middleware for space: robust asynchronous messaging; and a shared object model. We assume that communications services are progressively standardized (e.g., CCSDS protocols).

Robust messaging addresses many of the limitations imposed by the space domain, such as disconnectedness, long-latency links, etc. Moreover, message-oriented middleware (MOM) has a successful history in the evolution of publish/subscribe internet applications, as distinct from those built using tightly-coupled remote invocation. At a high level, this difference can be compared to that between sending email and having a phone conversation. The former allows the sender and receiver each to continue performing work independently of the (background) email transfer, while the latter requires both parties to be interacting simultaneously in real time (with some acknowledgment). The ubiquity and volume of today's email traffic attests to its utility as a productivity aid in the work environment. Moreover, in the space domain it is often impossible to provide real-time interactivity. Messaging can also assist automation, since tasks can be linked to the content of such messages. As a simple example, action might be triggered only when the value of a certain field received in a regular message reaches a certain range.

This leads to discussion of the second enabling technology: a shared object model can enable such messages to have significantly more meaning. An object of a particular type encapsulates specific properties and coded functions; if an object of this type is created by one application then a similar "proxy" object can be created and manipulated on a remote platform without transferring the entire object (which may be quite large). Separate application components (e.g., on different platforms) can thus interact efficiently because there are standardized rules for each to deal with the referenced objects. An example of this in the space domain is that a model of a sensor can be created on the ground and can contain functions like "take current reading" or "set calibration value". While the actual sensor may be on a remote spacecraft, ground software can interact with the local model of the sensor in exactly the same way as the spacecraft can, provided the ground and flight software share the object model. More importantly, these objects can be coupled by messages containing only object references and property values, thus providing an extremely efficient way to structure communications.

A simple example of such a service accessed through a standard interface is the CCSDS File Delivery Protocol (CFDP) [3], which enables efficient file transfer over such space links, and conceptually replaces the standard Internet File Transfer Protocol (FTP). Viewed from the perspective of enabling middleware, CFDP hides the constraints of the space link by appropriate use of buffering, accounting, and retransmission, while guaranteeing complete and ordered delivery of data in file units (as opposed to frames or a bit stream). More generally, it can be viewed as an internet "service" designed to succeed when standard Internet protocols fail because one or more of their design assumptions (like continuous connectivity) are violated.

At a higher application layer, middleware enables the deployment of an agent infrastructure, in which software modules perform specific roles in response to specific inputs. For example, a vehicle health monitoring agent could subscribe to various state values, and consider the implications of certain combinations of recognized faults).

4. PROBLEMS

While describing some of the potential benefits of the IPN, we have also mentioned several of the engineering and technological challenges, such as harsh environments and resource constraints. However, there are other non-technical issues that may pose more difficulty than the technical ones. These include cultural issues, such as finding a way to "negotiate" the use of shared infrastructure to the mutual benefit of all. The long history of isolated missions has not provided much opportunity for dealing with such issues. It is thus insufficient to provide a working solution to concepts such as coordinating multi-mission assets (e.g., via distributed planning) if the mission principals involved are unwilling to consider such an approach. Part of this approach should therefore be to define and implement a technology transition path. This should include process elements (e.g., how research and technology tasks are initiated, tracked and evaluated) and funding elements (e.g., who pays for what tasks and when transitions occur). A stretch goal would be to attempt to reduce the technology lifecycle for the IPN from >5 years to <2 years. Lessons learned from the rapid evolution of the terrestrial Internet could be particularly useful here, for example relating to the failure of high-tech companies who could not survive the concept of producing incremental product on a "web-year" lifecycle.

A major issue is thus to define what persistent capabilities could provide most benefit to most space applications at the lowest cost, and then to develop and deploy them in the optimal order. Rather than attempting to reach a broad consensus on this *before* starting down this path, an alternative approach is to be more pragmatic – namely, to quickly demonstrate some initial simulated capability to gain support for extending it progressively. In this area, the authors have so far addressed two prototype space applications. The first addresses the common difficulty of intermittent or pass-oriented communications, i.e., the desire to utilize more connection paths (even unplanned ones) in a seamless way. The second addresses the situation where sensor bandwidth significantly exceeds downlink bandwidth, thus suggesting semi-automated mechanisms to manage storage and prioritize downlink. Areas that have not yet been addressed via prototypes are reliability and availability, which become paramount for services in the space domain, and need to be considered much more carefully than is typical for internet-based services. For example, automated transparent service redundancy could be prototyped. Such demonstrations can also help determine when the service approach is likely to provide sufficient return on investment, and to quantify the risk/benefit analysis.

5. CONCLUSION

Space-based middleware can help shift focus away from the details of point-to-point remote communication and towards a high-level service architecture with increased capability for automation and cooperation among space assets. A first step is to define a service-oriented data architecture, high-level data objects, and mechanisms to allow efficient exchange between producers and consumers interested in particular attributes of such data objects. This also helps on-board applications to be insulated from inessential details (including the vagaries of the space communication). We believe that many of these details are better handled at the shared service layer (once for all), rather than at the individual application level (every time).

The IPN will offer equally exciting possibilities for the general public. High-end immersive environments could be made available to the public at facilities such as museums, with more modest capabilities available on the home desktop. The applications would be different – conceived to educate and to entertain. You could explore the remote environment – "What's over that ridge over there?" You could pretend you are the spacecraft – having its mobility and senses. You might be able to experience odors or weather conditions on another planet. See Figure 4. In this way, the IPN will "bring the planets to the public," and help inspire the next generation of explorers.

There is significant ongoing work in areas such as IPN architecture, space communication standards, protocol development, onboard autonomy, distributed science. Space middleware is intended to be seen as assisting this work. It is perceived that the larger problem to be solved is not technological in nature, but cultural.

Figure 4: IPN Public Use Scenario – *Experiencing the Remote Environment*

This work was performed at the Jet Propulsion Laboratory of the California Institute of Technology under NASA support.

6. REFERENCES

1. V. Cerf et al., "Delay-Tolerant Network Architecture": March 2003: http://www.dtnrg.org/specs/draft_irtf_dtnrg_arch_02.pdf
2. K. Fall, "A Delay-Tolerant Network Architecture for Challenged Internets", IRB-TR-03-003, Feb., 2003: http://www.dtnrg.org/papers/IRB-TR-03-003.pdf
3. K. Scott, S. Burleigh, "Bundle Protocol Specification", Mar 2003, http://www.dtnrg.org/specs/draft_irtf_dtnrg_bundle_spec_00.pdf
4. K. Delin, S. P. Jackson, "The JPL Sensor Webs Project: Fielded Technology", presented at the 1st internation Conference on Space Mission Challenges for IT, August 2003. http://sensorwebs.jpl.nasa.gov/resources/Delin_et_al_Fielded_Sensor_Webs_-_SMCIT03.pdf
5. N. Lamarra, "Space-Based Middleware", in IND News, May 2003

Norm Lamarra is a software developer in the Engineering and Communications Infrastructure Section of JPL's Information Systems Division. He holds B.Sc. and M.Sc. degrees from the University of Birmingham, UK, in Mathematical Physics and Electronic Engineering respectively, and a Ph.D. from UCLA in System Science (Communication Systems). He has been at JPL since 1994.