

Criteria for Evaluating Alternative Network and Link Layer Protocols for the NASA Constellation Program Communication Architecture

Daniel Benbenek¹ and Jason Soloff²
NASA Johnson Space Center, Houston, TX, 77058

Erica Lieb³
ITT Advanced Engineering and Sciences, Herndon, VA, 20170

Selecting a communications and network architecture for future manned space flight requires an evaluation of the varying goals and objectives of the program, development of communications and network architecture evaluation criteria, and assessment of critical architecture trades. This paper uses Cx Program proposed exploration activities as a guideline; lunar sortie, outpost, Mars, and flexible path options are described. A set of proposed communications network architecture criteria are proposed and described. They include: interoperability, security, reliability, and ease of automating topology changes. Finally a key set of architecture options are traded including (1) multiplexing data at a common network layer vs. at the data link layer, (2) implementing multiple network layers vs. a single network layer, and (3) the use of a particular network layer protocol, primarily IPv6 vs. Delay Tolerant Networking (DTN). In summary, the protocol options are evaluated against the proposed exploration activities and their relative performance with respect to the criteria are assessed. An architectural approach which includes (a) the capability of multiplexing at both the network layer and the data link layer and (b) a single network layer for operations at each program phase, as these solutions are best suited to respond to the widest array of program needs and meet each of the evaluation criteria.

I. Introduction

SPACEFLIGHT presents challenges for the design of communications and network infrastructure. Communication designs for spacecraft must contend with the dynamics of spaceflight and must include avionics suitable for the space environment. Manned space flight (MSF) presents additional challenges beyond unmanned flight, with a need for increased levels of connectivity and reliability, supporting real-time operations and crew safety. In the selection of a communications and network architecture, an emphasis should be placed on the flexibility to accommodate changing missions and mission objectives. The necessary applications, such as voice, motion imagery, file transfer, telemetry and command, will remain relatively constant across time and mission

¹ Computer Network Engineer, Constellation Program Office, NASA Johnson Space Center, Houston, TX 77058 . .

² Avionics and Communication Lead, Constellation Program Office, NASA Johnson Space Center, Houston, TX 77058

³ Systems Engineer, ITT Information Systems, 12975 Worldgate Dr. Herndon VA 20170, AIAA Member

objectives; it is the network and communications functions (network and data link layer) that will determine the capability of the end-to-end architecture.

An evaluation methodology is required to determine if potential communications and network architectures will meet particular program needs as they evolve. In Section II an introductory discussion of possible program goals and objectives, spanning from lunar sortie to Mars and flexible path options for travel to near-Earth objects (NEOs) provides an overview of the communications and network needs and challenges. Section III introduces and defines criteria against which potential architectures may be compared, and Section IV follows with an assessment of three key architecture trades.

The criteria for evaluating a potential architecture against a given use case are essentially the same drivers for any network decision, but guided by the unique needs of MSF. These criteria are not expected to be static for the life of the program, but currently include interoperability, security, reliability, and the relative ease of automation of topology changes.

Among the architecture options that can be evaluated, three key trades are highlighted: (1) multiplexing data at a common network layer vs. at the data link layer, (2) implementing multiple network layers vs. a single network layer, and (3) the use of a particular network layer protocol, primarily IPv6 vs. Delay Tolerant Networking (DTN).

This paper concludes in Section V by proposing a recommended architecture based on the considerations of the evaluation criteria, architecture trades, and current program objectives.

II. Exploration Use Cases and Their Communications and Network Needs and Challenges

Consistent with the Vision for Space Exploration (VSE) (Ref 1) established in 2004, Constellation Program objectives are currently focused on initial lunar return missions (sorties) progressing to a sustained presence (“outpost”) and eventually on to Mars. These objectives carry with them an implied evolution in capability, including increasing mission autonomy and less reliance on real-time Earth (Mission Control) involvement. The use cases for the trade space described in this paper are based on current program objectives as well as possible departures and various evolutionary paths. The selection of use cases and the sequence in which they are pursued will be the drivers for the selection of a communications and network architecture. Some key discriminators include scope of mission (number of participating assets), diversity in location, reliance on Earth or tolerance to Earth communications delay, and the degree to which the local network must be self-sufficient and support more autonomous operations.

An additional driver that will determine architecture is the tradeoff between simplicity and flexibility. In the conservative world of manned space flight, capability above the minimum required to accomplish the requirements, even at no extra development cost, is not always considered a positive. As with any man rated system, unpredictability of any part of the spacecraft is a safety concern. To minimize unpredictability crew and mission support are trained, to the extent feasible, for all situations, and likewise, to the extent feasible, all configurations are tested. In this way a more flexible system may incur greater cost to implement and maintain, even if direct development costs are minimal.

A. Lunar Sortie

Prior to establishing a sustained human presence on the moon, a number of initial robotic and human lunar return missions will be conducted. These missions may explore a variety of different locations, scouting for the best possible outpost location, conducting unique science, and preparing for a more sustained presence.

From a communications and network architecture perspective the sortie class missions have unique needs. The architecture should flexibly support diverse locations, provide data and imagery transport that enables command and control of initial robotic missions from Earth-based ground stations, and be geared toward focused support of a relatively small set of assets operating on the surface (e.g. lander, EVA crew, and rovers). For the first return missions to the moon, public outreach will be important and the architecture should support high definition still and motion imagery in near real-time. Operationally, there is a greater likelihood that early missions back to the moon – both robotic and crewed – will rely heavily on Mission Control presence and direction, driving the requirement for near continuous real-time communications and networking between the moon and Earth.

B. Lunar Outpost

As our permanent presence on the moon comes into being, the lunar outpost site will gain assets with each landing, and the activities will become more complex and diverse; roving science assets, in-situ resource utilization

efforts, assembly and additions to habitation by robotic and EVA (or IVA) crew, power generation and distribution, and so on. A single sustained outpost may allow for a static or single-location network, but the flexibility to conduct sortie explorations in conjunction with the outpost activities would demand that the architecture continue to support diverse locations.

The communications and network architecture must arguably be designed to evolve from the requirements of a sortie mission into a more robust and extensive network providing two key characteristics: (1) data transport for an increasing number of surface elements; the network should support expansion in the number of network nodes or participants in the network, and (2) equal if not greater emphasis on robust local (surface-to-surface) communications and networking allowing the crew at the outpost greater visibility into and control of their local environment. Although Mission Control would continue to be a presence in lunar outpost operations, the outpost would begin operating with a larger degree of autonomy; continuous real-time communications with Earth may be available but used to a lesser extent. This operational transition allows simulation of Mars-forward activities, allowing crew and Mission Control to operate as if a Mars light time delay were in effect. Implications of a sustained presence and operating with greater autonomy include increased emphasis on the reliability of local networks (inclusive of hardware and software). The length of each mission and the need for self sufficiency will mean that the network should be relatively robust and allow for upgrades or replacement of elements if necessary.

C. Mars

What will be simulated in later years on the moon, will be reality on Mars. The most significant paradigm shift for the communications and network architecture will be the impact of light time delay; one way light time delay varies from 3 to 22 minutes, eliminating the feasibility of near-real time control of surface elements from Earth as well as limiting if not completely eliminating the possibility of Mission Control actively engaging in resolution of an emergency. The architecture must accommodate the impacts and limitations of the communications delay, and similar to the lunar outpost, the network's ability to enable the crew to be self-reliant and maintain situational awareness of their local environment will be critical. A key to enabling self-reliance in this remote, unforgiving environment will be an infrastructure that does not burden the crew. Specifically, the network must be automated in establishing and maintaining necessary communication paths and must provide security with minimal user impact.

D. Alternative Options and Evolutionary Paths

In light of the Review of U.S. Human Spaceflight Plans Committee Report (Ref 2) which included a "Flexible Path" option in addition to Moon and Mars destinations for exploration, it is worth noting what the communications network architecture drivers might be for such an option. Included in the notion of the Flexible Path option are inner solar system locations such as lunar *orbit*, Lagrange points, near-Earth objects (such as asteroids and comets) and the moons of Mars, all as potential precursors to lunar and/or Mars surface destinations. For the most part, these flexible path options would have some similarity to the lunar sortie mission with respect to the relative size of the mission and number of assets involved and participating in the mission's communications and network architecture. The variety in location spans the modes and characteristics highlighted in the lunar and Mars use cases; in some options a near-real time network connection with Earth may be available, and in others the more self reliant and robust local network may be required and communications with Earth must be adaptable to the light time delay.

III. Communications and Network Evaluation Criteria

The criteria for evaluating a potential architecture against a given use case are essentially the same drivers for any network decision, but guided by the unique needs of MSF; increased levels of connectivity and reliability, supporting real-time operations and crew safety. The criteria characterized below include interoperability, security, reliability, and the relative ease of automation of topology changes, which includes mobility capabilities.

A. Product Interoperability

The notion of interoperability continues to gain traction as increasing emphasis is placed on international cooperation and involvement of commercial partners in future exploration initiatives. Two key activities defined in the VSE as supporting or enabling the established goals were: (1) Pursue opportunities for international participation to support U.S. space exploration goals, and (2) Pursue commercial opportunities for providing transportation and other services supporting the International Space Station and exploration missions beyond low Earth orbit. Within

NASA's Space Operations Mission Directorate, the Space Communication and Navigation group is investigating commercial interest in development of future lunar networks (Ref 3) and has included commercial standards in concept communications and navigation architectures (Ref 4).

The Inter-agency Operations Advisory Group, which provides a forum for identifying common needs across multiple international agencies for coordinating aspects of space communications, chartered the Space Internetworking Strategy Group (SISG) in June of 2007. With a charter to establish an approach for transitioning toward network centric space mission operations, defining and establishing options for interoperability has been a key activity of the SISG. The SISG has adopted the following definition of interoperability: the technical capability of two or more systems or components to exchange information and to use the information that has been exchanged (Ref 5).

For purposes of this discussion an expanded definition is proposed under the heading "product interoperability." The ability for two or more systems to exchange and use information may be achieved on a number of levels. One might think of engineering interoperability as the capability to work and coordinate amongst partners to make a solution workable; this is the Apollo 13 model, where the square-peg-in-a-round-hole dilemma is solved via plastic bags, cardboard, and tape. Product interoperability, on the other hand might be thought of as designing square pegs for square holes from the start; an interoperable architecture on the product level would enable a truly "plug and play" environment.

Characteristics of an architecture designed with product interoperability include: common avionics that are easily replaceable or exchanged between systems of the same program or between partners, resilience to changes in vendors, and cost benefits associated with a common product line.

Interoperability is more difficult to achieve in the manned space flight environment, because of the minimalistic approach to software design. Options and redundant methods tend to be avoided in implementations, even if specified in standards.

Interoperability can be required at any layer in the communications stack. At the physical layer, interoperability is generally available, because programs/users tend to use well known industry standards without much customization.

Issues begin to arise in the data link layer. As a result of the negotiation and compromises associated with the standards approval process, the specifications are often written such that more than one implementation approach will meet the standard. When spacecraft communications are designed, almost invariably, the minimalistic approach to software results in a tailoring out of all specification alternative methods, save one. Implementations of the same data link layer protocol for different projects will be divergent and tend to be non-interoperable. We see this with implementations of the CCSDS Advanced Orbital Systems (AOS) specifications.

Similar issues are arising as the space programs begin to add routable network layers to the protocol stacks of their flight systems. A specification considered mandatory in an Internet Engineering Task Force (IETF) Request for Comment (RFC) may be excluded if the designers can reasonably conclude that it is not needed. Again divergent implementations will take hold, each of which will be derived from "the standards" but which may be non-interoperable.

B. Security

Security in a communications and network architecture is the protection of information that is created, used, stored, and/or communicated in support of the mission(s). The need for security is driven by regulation, safety, and privacy concerns. Although some regulations stipulate protection of data records, for purposes of evaluating communications and network architectures the focus is on protection of the process (data transmission).

Traditionally, the concepts of data integrity and, increasingly, authentication, have been the most important foci of computer security for space flight. The responsibility for data integrity tends to be shared throughout the stack as its assurance is necessary for the continued functioning of the communication flow. Checksums or cyclic redundancy checks (CRCs) take place at many layers and are available for most protocols, such that data integrity will generally not represent a discriminating factor for choosing architecture.

Command authentication is mandatory for assets of any significant value. The value of return link authentication, relative to cost, for a small project is debatable, and past civilian space programs have tended to omit it. The reasoning is that the trouble and expense of fabricating a downlink signal for the purpose of impersonation would prevent its occurrence, while frequency allocation and the custom nature of higher layer software prevents accidental reception or processing of the signal from another vehicle. For large manned space flight programs, such as Constellation, these arguments may not hold.

With command and control paradigms shifting from commanding through a single control center to an eventual peer based architecture in which space based assets will send commands to each other, authorization becomes an important information security topic. Even within a trusted network, configurable validation of the command source will be necessary to prevent overlapping or errant commanding.

In manned space flight (MSF), confidentiality mechanisms are needed to protect the privacy of the crew. Any data related to the health of the crew or crew support will, at the very least, raise questions of confidentiality.

Mechanisms for meeting the security needs of authentication, authorization, or confidentiality occur at various layers on the protocol stack. This paper will examine the relative merits of link layer security (LLS), network layer security (NLS), and application layer security in meeting the identified needs.

LLS almost always takes the form of blanket authenticated unidirectional encryption of everything past the data link header. NLS also tends to be used to provide blanket coverage, but can be used selectively to protect some data while allowing other to flow freely. Application layer security moves the burdens of security to the end points, which can be problematic if the goal is to build a secure network.

C. Reliability

Reliability can be defined as the probability that a device, product, or system will not fail for a given period of time under specified operating conditions (Ref 6). Reliability is a characteristic of the architecture design, impacts operations and maintenance costs, and may impact loss of mission and loss of crew members. As such, reliability may be seen as investment in risk reduction at the price of more robust or redundant implementation choices.

A reliable communications and network architecture will have a mean time between failure (MTBF) equivalent to multiples of the mission duration, and will exhibit fault tolerance or graceful degradation in the presence of failure. Fault tolerance is not limited only to individual systems but may be a characteristic of how systems interact together. Often the ability to avoid completely any kind of failure is simply too cost prohibitive. Instead, an approach that minimizes the impact of the failure, gracefully degrading services in proportion to the severity of the failure rather than terminating them, may be optimal.

In the quest to maximize “uptime” the role of protocols that provide a self-healing capability and a topology designed to leverage that capability comes into discussion. Routing protocols are an obvious example of this self-healing capability and also can illustrate the limitations of such an approach. The devices running a routing protocol require additional configuration, and yet the routing protocol will provide limited benefit on a network that only has one path available.

D. Ease of Automation of Topology Changes (EATC)

The ease with which network topology changes can be made in an autonomous fashion may be considered a multi-part architecture attribute. EATC includes the ability to adopt new systems and technologies, adapt to increasing levels of system autonomy, allow for robust fault detection and resolution (self-healing), and to detect and utilize network nodes and paths in a dynamic environment. The nature of exploration requires this kind of capability; the fact that our lunar, Mars, and/or “flexible path” exploration will be a multi-decade adventure drives the architecture to incorporate advances in technologies and standards, as well as the introduction of new elements into the network. In the context of a lunar outpost the need for sensing and engaging new network nodes is evident. The number of surface elements will continue to grow (mobile habitats, rovers, science assets, power plants, ISRU deployments, robotic assistants, and so on) and as the elements move in/out and around the outpost zone the connectivity between elements will change accordingly.

Changes occurring at the physical layer (e.g., an RF link drops due to range), can lead to changes at the application layer. For example, an application may write to memory when the network is down and stream data when a remote location is reachable.

EATC is supported by modularity, which will be a cornerstone of any sustained exploration program that must develop and deploy assets over an extended period of time. More critically, an architecture that incorporates EATC assumes that systems will improve over time and will continue on a course toward autonomy. The value that an MSF program places on EATC reflects a program’s acceptance of both the concept that autonomy from Earth based control will increase over time and the concept that increasingly sophisticated software is one key to increasing autonomy. Since the software providing EATC will be introduced into an existing network infrastructure, a built-in modularity to the software structure will greatly aid in realizing the benefits of EATC.

The network architecture must accommodate two types of change in topology. The first is a transient topology change. End nodes may come into and out of the network if they are mobile or if they are simply not powered regularly. This change can occur at the end node level or at the core of a network. An EVA suit is an example of an

end node which may be added to the network as needed. An orbiting satellite is an example of a network core transient body. When in sight of both surface (lunar or Mars) and Earth resources the satellite can act as relay. As surface operations become more complex, accommodating transient topology changes through automation will become increasingly important.

The second type of topology change to be accommodated addresses an ability to alter peering relationships in a dynamic network. A four node network may form linearly such it takes three hops to forward data from node A to node D. Alternatively, all nodes may be peers, each with the ability to serve as relay for other nodes. In such a topology, nodes A and D are simultaneously one, two, and three hops apart and one should rely on the automated protocols to pick the best path, which may or may not be the most direct.

The following are some examples of software supporting EATC that could be important to manned space flight. On terrestrial networks, the availability of physical links can trigger an automatic introduction of data link layer and network layer. A similar software hook between RF based physical links and the corresponding data link layer will maximize network availability without burdening human network operators. Mobile-IP provides for automated changes in physical topology while retaining the apparent location of nodes in the network layer. Sophisticated, but low overhead mesh link layer and corresponding network routing protocols will allow end nodes to move from hub and spoke to daisy chain topologies transparently to the users. Finally, DTN is predicated on network layer awareness of link availability to relay data without the necessity of operator cognizance.

The upside of EATC is that operator interaction with the network is generally minimized. The downside of EATC is that software which functions automatically often requires more upfront development and configuration and that troubleshooting often requires a higher degree of expertise.

IV. Communications and Network Architecture Trades

There are many tradeoffs that must be made in constructing an end-to-end network. Among the first and most basic decisions is the determination of the data link and network layer protocols. For the purposes of this paper we have expressed this determination in terms of three specific trades: (1) multiplexing data at a common network layer vs. at the data link layer, (2) implementing multiple network layers vs. a single network layer, and (3) the use of a particular network layer protocol, primarily IPv6 vs. Delay Tolerant Networking (DTN). Advantages and disadvantages of the approaches are described and evaluated in terms of the evaluation criteria outlined in Section III.

A. Multiplexing Data at the Common Network Layer vs. at the Data Link Layer

Multiplexing of application data is one of the most important functions of modern networks. Practically speaking, data is either multiplexed at a network layer or at the data link layer. Multiplexing at the transport layer is considered a subset of network layer multiplexing, because for all practical architectures the available transport layers are tied to the network layer. Multiple data link layers are not run simultaneously over the same physical interface. Multiplexing above the transport layer is sometimes efficient for individual services (e.g. voice bridging, command consolidation) but is not practical for combining multiple disparate services.

On ground networks, multiplexing at the data link layer would generally mean running multiple network protocols (for instance, running IPv4 and IPv6 over the same physical installation). For MSF, this paradigm applies to internal system networks and intersystem links that are over umbilical as opposed to RF.

For RF links, NASA and other space agencies have committed to using CCSDS standards and protocols. Furthermore, NASA MSF has restricted itself to AOS as the data link layer over RF links. Use of a single data link layer for all RF would simplify software design for multiplexing operations at the data link layer. For our trades, we assume that multiplexing at the data link layer means that AOS is the data link layer for all RF, and that Ethernet is the data link layer for all internal networks and intersystem umbilical connections. More complicated scenarios are possible, but the very complexity of those scenarios is likely to reduce the attractiveness of multiplexing at this layer.

When considering multiplexing at the network layer, both IP and DTN will be discussed. The focus will be on similarities, while other sections will emphasize differences between the two protocol families.

Whenever interoperability is a prime emphasis, a single network layer will most likely be the resulting architecture, unless the topology of the network is very simple, without the need for systems to forward data for each other. Use of a single network layer tends to point towards multiplexing only at the network layer. If one wants all systems to be reachable by all other systems, then intermediate systems must forward data for end nodes and for

each other. Separation of data into virtual channels and virtual LANs works against the objective of simplifying data forwarding.

Multiplexing at the data link layer allows for segregation of traffic in separate virtual networks. Use of different network layers (or absence of a network layer) is one reason to support virtual networks, but not the only one. One might also consider a limited multiplexing at the data link layer as form of pseudo “out-of-band” communication, preserving a communication path in case of network layer anomalies or heavy utilization on the primary channel. Use of different security levels has also been used as a reason to keep data on different virtual networks.

B. Single vs. Multiple Network Layer Architecture

When all nodes on a network interface with the same network layer, it means that only complimentary application layer software needs to be added to enable any two nodes to interoperate with each other. In the case that interoperability for all nodes is perceived to have a high value, use a single network layer is generally taken as a given.

Using multiple network layers provides flexibility for accommodation of heterogeneous systems. This approach would generally be favored in the case of multiple independent parties making use of the same infrastructure. An infrastructure designed to support for multiple network layers will also allow for gradual upgrades as opposed to across-the-board uniform upgrades. While security is not necessarily enhanced by the use of multiple network layers, multiple network layers permit greater variation in security approaches used by the entities using the network. In short, autonomy is maximized when multiple network layers are permitted at the cost of interoperability and ease of automation.

C. Internet Protocol Version 6 vs. Delay Tolerant Networking

The importance of choosing a network layer protocol is obvious if the choice has been made to use a single network layer architecture. However, even if multiple network layers will be supported, each network partner, administering an independent network layer architecture over the physical infrastructure, must make this evaluation.

Many of the touted advantages of IP are those that properly belong to the network layer in general. Likewise, many of the acclaimed benefits of IPv6 are given in comparison to IPv4. The problem with the IPv6 vs. v4 comparison is that inevitably it becomes a comparison between current vs. planned capability. Without getting into a debate of the relative merits of v4 vs. v6, we would make a claim that the IP of the future is an evolution from the current IPv4 suite using the IPv6 address model with additional IPv6 features added as needed .

An example used as evidence of this evolution is the use of IPsec within IPv4. IPsec was originally conceived as an integral part of IPv6. The reality, however, is that IPsec changes, fixes, and enhancements take place first for IPv4, because that’s where the market is.

Assuming that the industry trend of evolution vs. revolution for network layer communication to be self-sustaining, our layer three comparison is between DTN and IP, which, to avoid the argument begun above, we will say is identical to IPv6. However, one runs into similar sorts of issues when attempting to compare DTN to IP. With DTN still in the developmental phase, one cannot state for certain, a capability that it will NOT have, particularly for our direct comparison. Furthermore, if the applicable timeframe of DTN to MSF continues to move to the right, then it is conceivable that commercial IP suites may incorporate many DTN features, as well.

However, ignoring the speculative possibility of IP and DTN convergence into a single multifaceted network protocol, there are still discriminators to be considered. IP is the dominant network layer protocol on the planet. Expertise in IP is common, wide-spread, and relatively cheap to employ. IP software code has been well-tested for interoperability and debugged in the market place. On the downside, IP code also tends to contain obsolete or archaic branches which are no longer needed for real-world internet communication let alone MSF.

By contrast, DTN code has the possibility of being more streamlined. With the inherent dangers in spaceflight, NASA tends to develop custom code with which developers and testers become intimately familiar. However, assuming DTN does not become a widely-used protocol suite, the concept of leveraging widespread expertise would not apply.

If code is highly customized then no significant advantage is gained from a wide engineering base, which tends to minimize the advantage of IP over DTN. Mitigating this last point, however, is the concept that, for non-critical service, commercial applications could utilize the same network layer infrastructure as critical NASA-developed IP apps, while new development would be needed for each DTN corollary service.

DTN is often described an overlay network (Ref 7), providing an end-to-end path for communicating nodes through a heterogeneous infrastructure which would not otherwise support the required connectivity. However, with the minimalist approach favored by flight software developers, NASA is unlikely to develop layers for a

communication stack if one can provide the required functionality. Informally, it has been suggested that DTN alone can provide all necessary network layer connectivity for at least the space portion of exploration connectivity.

DTN is actually a catch-all description of a number of efforts, often overlapping, within DoD, IETF, CCSDS and others to develop layered communication software tolerant of interrupted service. Any MSF or other major NASA program implementations of DTN would almost surely be drawn from the CCSDS DTN standards. If one accepts an IP plus model of DTN, it becomes difficult to argue against DTN. The remaining legitimate arguments for comparison are the speed of development and maturation of DTN vs. the timeframe need of the program in question, and the virtues vs. vices of using well-known, established protocols and code.

V. Selecting a Communications and Network Architecture for Exploration

By establishing evaluation criteria (established in Section III) and defining the major trade space (Section IV) proposed architectures for different missions and mission needs can be assessed in methodical fashion.

A notional mission manifest is proposed for illustration and discussion of the architecture options in the following subsections:

- LEO operations/ISS docking 2014
- Lunar Sortie – Initial Human Lunar Return 2020
- Lunar Outpost 2025
- Mars 2040

A. LEO Ops and Lunar Sortie

The start dates for LEO operations and the first lunar sortie missions are important as they play a role in determining protocol choices based upon assumed maturity level. For both LEO ops and initial lunar nearside sorties, where human spaceflight is concerned, communication will be nearly constant and the diameter of the network (the maximum number of hops) will be relatively small. Simplicity and reliability will be the most important drivers. Cost will be crucial, as it is with all government funded programs. The desire to control costs will tend to limit changes to those of an evolutionary nature. Throughout this period, MSF will be learning to operate in a network layer based communications infrastructure. The experience gained will be crucial to administering the more complex network topologies which will emerge in the later Program phases.

In terms of multiplexing at the link layer vs. network layer, outside of preparation for future phases, neither presents significant advantages over the other. Because of the small size of the network a single network layer will be assumed. Exceptions may occur when testing a planned change or expansion in network protocol. Specifically, the Cx program will start with IPv4, and will likely eventually migrate to IPv6 or DTN. Early testing of a new network protocol in flight would amount to support of two or more independent networks multiplexed at the link layer. However, the topologies of both network layers will likely be quite simple so as to create few additional challenges. Whether the additional layer (or layers) is IPv6, DTN, or both, depends on the architecture chosen for the later phases. Lunar far side sortie does not represent a separate program phase; rather it is mission scenario which, when present, will drive the capabilities of the communication system, as well as many other subsystems towards autonomy from Earth based operations.

Let us assume the following systems for a far side lunar sortie: Two Extra Vehicular Activity (EVA) suits, other end systems such as personal computing devices, a rover, a lander, an orbiting relay satellite, and, of course, ground stations and mission control on Earth. Let us assume that excursions via rover are part of the mission plan, and that constant contact with mission control is not considered mandatory even during excursions. Such assumptions indicate the need for a network topology that is self-forming, self-healing, and disruption tolerant. The network is not sufficiently complex as to warrant parallel networks, except possibly in a case of an emergency communication channel.

If one accepts that eventually DTN is necessary for solar system exploration (at Mars distances, for example) then lunar far side sortie becomes prominent in driving network architecture, as the first scenario for which disruption tolerance is clearly useful – highly desirable, if not absolutely mandatory. This is true to such a degree that if far side sorties are manifested early in the lunar phase, that it drives the architecture to early adoption of DTN. The size of the overall network during sorties is not sufficiently complex to justify running both IP and DTN in parallel networks. This suggests a single network layer protocol, DTN, towards which all applications are multiplexed, except perhaps for emergency comm.

For emergency communications, let us consider a driving EVA scenario, typically referred to as a “walk back” scenario. In the walk back scenario, EVA crew on a roving excursion are forced to return to a safe haven on foot due to a vehicle failure; in early planning discussions within CxP operational guidance has been limiting the walk back distance to 10km. For a walk back scenario during a far side sortie with a rover excursion, the communication path will be EVA suit to relay satellite to Earth (ground station and mission control center). We consider two options for the emergency communication architecture to accommodate walk back: voice-only and low-bandwidth DTN. A voice-only link would consume the lowest power, but it might be more costly from an EVA standpoint in terms additional electronics and complexity. A low bandwidth DTN solution from the EVA suit directly to the relay satellite allows mission control to assist the astronauts during the walk back both through diagnosis of the suit and through voice communication which could function over an intermittent end-to-end path. Survivability then becomes the issue. If walk back requires power conservation, then disrupted communication may be considered acceptable. If a low bandwidth DTN solution is deemed acceptable from a power conservation standpoint, then the more consistent and functional communication platform would most likely enhance survivability.

Other scenarios besides walk back may drive a need for out-of-band emergency communications. If out-of-band emergency communications is deemed necessary, then multiplexing at the link layer is indicated. This then becomes an assumption for the lunar architecture in general, and we would assume that multiplexing will occur at both the network layer and the link layer.

B. Lunar Outpost

If a lunar outpost is established it will likely occur through a gradual growth in deployed equipment and overall capabilities with each mission. The overall complexity in terms of deployed assets will drive the resulting network architecture in terms of multiplexing and network layer choices. So, in addition to the scheduled appearance of far side sorties, the final size and complexity of the network during the lunar phase will drive overall lunar architecture choices.

For a view of the greatest complexity to be reached for the lunar phase network, we start with the assumption that International Partners will play an important role. Although we hope to build on ISS lessons learned, it is [also](#) reasonable to assume that some of the precedents set in ISS will likely continue, namely that international partners will desire independent control of their resources but a highly reliable primary communications path will be shared.

We see the optimum situation as one in which all systems and partners run the same network layer protocol. This will maximize interoperability and versatility of communications relays and even endpoints, enhancing survivability in the same way that hardware standardization can. This does not automatically imply that all communication will occur over the same virtual network. Communication using the same network layer may be parallel and independent, by multiplexing at the link layer. Link layer multiplexing would not be as efficient in terms of bandwidth utilization as multiplexing at a common network layer using traffic prioritization, but it may be necessary in order to minimize disagreements in bandwidth allocation. The desire by partners to maintain network independence should not be a reason to maintain different network protocols.

However, budget constraints will always be used as a justification for continuation of existing architectural stovepipes. Some lunar architecture discussions have portrayed IPv6, DTN, and Space Packet all flowing over the same AOS infrastructure. How data forwarding would occur in such architecture is unclear. Separate routers for each network protocol or custom routers supporting multiple network protocols will be more costly and of lower reliability than if the architecture were to use a single network layer.

There should be a list developed of IP-like qualities that must be included in the DTN suite in order for Cx to consider replacing IPv6 with DTN in the lunar time frame rather than the current Cx plan to overlay DTN on top of IP. The following is offered as a starting non-comprehensive list:

- Integrated security as robust as that provided by IPsec.
- A name resolution and addressing scheme that supports isolated or universal connectivity
- Optional name/address automation to support automated connectivity
- Support of acknowledgement free data streams (UDP-like)
- Support for data streams which require real-time acknowledgment (TCP-like)
- No impediments to using both ack/non-ack simultaneously from same node even same application (RTP/RTCP-like)
- Support for application/source/destination filtering equivalent to an IP firewall.
- Support for mobility (e.g. mobile IP)

- Support for multiplexing at the network layer without additional software: There are multiple IP types in addition to differentiation at the transport layer (tcp/udp ports).
- Routers that nominally support various data flow types (ack, non-ack, bundles) simultaneously

If the lunar outpost is dropped from the long term exploration plan then the level of communication complexity required to support for what will be primarily NASA sorties goes down. If far side sorties are not part of the lunar phase plan or only occur at the end of the lunar phase, then the benefits for Cx to migrate from IP based communication to DTN in the lunar phase are limited. One other scenario needs to be considered as a driver for network architecture in the lunar time frame. It may be desirable to simulate Mars-like connectivity conditions while maintaining primary communication links. A lunar outpost may support parallel network layer infrastructures, IP for normal ops, and DTN for Mars environment testing. This could be a precursor to total migration to DTN even if the migration takes place in the lunar timeframe. Depending on the complexity of network being simulated, one might encapsulate DTN over an IP infrastructure or carry both DTN and IP on separate virtual channels to Earth ground stations.

C. Mars Missions

Most analysts agree some sort of disruption tolerant communication architecture will be needed in order to support disjunctive communication between Earth and a Mars bound spacecraft carrying humans. We take as a given that DTN will be used to communicate from the Martian environment back to Earth. Hopefully international cooperation will have evolved to the point that parallel infrastructures are unnecessary. Inherent in this latter assumption is that network and application security is implemented where necessary to support privacy, authentication, and authorization needs. The remaining question is what sort of communication will be in placed between systems in the Mars sphere of influence.

The Mars area communications architecture will be dependent on the complexity of the missions and number of assets involved. Assets on opposite side of the Martian globe will have communication challenges similar to communication between Earth and the far side of the Moon. Such a case would point towards standardization on DTN for all network layer communications. In we make the assumption that Cx will want to maximize flexibility by allowing crew to do their own mission planning, once again, a DTN only architecture makes sense. This, of course, assumes that DTN lives up to the promise of incorporating most features of IP.

Multiplexing at the link layer may be preserved for emergency comm, as discussed above. In this case the emergency response would be that which is enabled by the infrastructure on and around Mars.

D. Other Missions (Asteroid, Comet, Lagrange Points, Mars' Moons)

None of these scenarios are currently part of the Cx program plan. As such, the goals and by extension, the network performance requirements for such mission will be highly speculative. However, two discriminating factors should be influential in analyzing network architecture to support any of these types of mission: Distance from Earth and fly-by vs. touchdown missions.

Fly-bys will have simple network topologies very similar to the LEO scenarios that do not involve docking. If all such activities took place sufficiently close to Earth that real-time communication can remain interactive and constant, then any changes in network capability would likely be implemented solely to prove out concepts for future mission.

Sorties in which a human occupied lander is involved should have many similarities with lunar sorties, with rovers and EVA suits included in the overall communications architecture. However, for scenarios relative close to Earth the topological variations of the network should be relatively small. Again one might refrain from adding the capabilities needed for more complex scenarios unless, again, as a proving ground for future missions.

When missions extend beyond a certain distance from Earth then real-time interactive communication is no longer practical. The delay tolerant aspects of DTN begin to be needed even if communication is continuous. Missions to small bodies at Mars or greater distances will require the similar levels of autonomy as actual Mars missions, but again we would assume that the complexity and potential topology variations of the network would be much less.

Manifest flexibility is another potential factor that may influence the timing of network architectural changes. If all immediately identified missions are relatively short distances from Earth, but the ability to reach further destinations must be supported in less time than required for major communication redesign, then the network capabilities must be included at the start of the program phase that includes these "other missions."

An important part of the Cx C3I vision has been an assumption that the exploration program can make upgrades to its communication and software capability. The move from IPv4 for ISS Ops, to IPv6 for lunar and DTN for Mars reflects this optimism. Some experienced MSF engineers remain skeptical regarding the ability to implement major changes in software once a program is operational. Recognizing in advance, the realistic opportunities for upgrade and planning to take advantage of those opportunities will be an important role for communications engineers of the future exploration programs.

If disruption tolerance capabilities are incorporated into the mainstream commercially market of Internet Protocol Suite (IPS) implementations, then the planned program upgrades may take the form of realignment with what would then be current Internet capabilities. Such an approach would put network layer communications on a similar track with operating system capabilities.

VI. Conclusion

The primary drivers for the network architecture for the various phases of Cx will be the manifest, the complexity of the required communications at each point in the program, and the intended autonomy of the deployed resources, which may be closely related to the distance from Earth. The timing of lunar far side sorties and the establishment of a lunar outpost will be particularly important in timing a migration from the IPv4 based system of the initial capability to a likely DTN centered architecture for Mars missions.

Use of a single network layer promotes interoperability, eases the implementation of automation, has little to no downside from security standpoint, and will likely be more reliable due to simplicity.

Multiplexing primarily at the network layer allows for the most efficient use of bandwidth, while preserving the capability to multiplex at the link layer aids in some contingency and makes traffic separation easier to manage.

With the lunar mission outline appearing to contain the most significant variables, interagency work to develop shared standards should continue in earnest. Any significant delay in lunar missions may allow for convergence on network layer standards, which would indicate a single network layer for the Cx lunar phase.

Acknowledgments

This paper was developed by members of the Constellation Systems Engineering and Integration, Software and Avionics team. The authors wish to thank the broader team and Constellation community; it is our hope that lessons learned from the Constellation Program to date support the next steps in manned space flight and exploration. Special contributions from Larry Schneider / COLSA, William Walsh / Linquest Corp., Jay Gao / JPL are appreciated.

References

- ¹ NASA, "The Vision For Space Exploration," NP-2004-01-334-HQ, February 2004.
- ² Review of Human Spaceflight Plans Committee, "Seeking a Human Spaceflight Program Worth of a Great Nation," 2009.
- ³ Commercial Lunar Communications and Navigation Request for Study Information, August 15, 2008.
- ⁴ Schier, James, "NASA's Lunar Space Communication and Navigation Architecture," AIAA Paper 092407.
- ⁵ Report of the Inter-agency Operations Advisory Group, Space Internetworking Strategy Group, "Recommendations on a Strategy for Space Internetworking," July 28, 2008.
- ⁶ NASA Systems Engineering Handbook, Rev1, NASA/SP-2007-6105, December 2007.
- ⁷ Fall, K., "A Delay-Tolerant Network Architecture for Challenged Internets," SIGCOMM, August 2003.