

Networked Operations of Hybrid Radio Optical Communications Satellites

Alan Hylton
alan.g.hylton@nasa.gov¹

Dr. Daniel Raible
daniel.e.raible@nasa.gov²

In order to address the increasing communications needs of modern equipment in space, and to address the increasing number of objects in space, NASA is demonstrating the potential capability of optical communications for both deep space and near-Earth applications. The Integrated Radio Optical Communications (iROC) is a hybrid communications system that capitalizes on the best of both the optical and RF domains while using each technology to compensate for the other's shortcomings. Specifically, the data rates of the optical links can be higher than their RF counterparts, whereas the RF links have greater link availability. The focus of this paper is twofold: to consider the operations of one or more iROC nodes from a networking point of view, and to suggest specific areas of research to further the field. We consider the utility of Disruption Tolerant Networking (DTN) and the Virtual Mission Operation Center (VMOC) model.

I. Introduction

The goal of this paper is to both provide an anchor point for the state of networking research as applied to the Integrated Radio Optical Communications (iROC) project and to list several research projects needed by iROC and the general space community.

All optical missions, NASA's or otherwise, are hybrid radio optical missions - all optical missions have radio support. The iROC project takes this a step further by considering the radio and optical links as members of the same system. This is achieved in hardware with a hybrid telescope-antenna design, and in software via the Disruption Tolerant Networking (DTN)^a protocol.

DTN¹ is a store and forward network overlay that can operate over heterogeneous subnetworks. DTN provides autonomous link management, buffer management, security for applications. DTN also has a quality of service (QoS) mechanisms to prioritize and offers a standardized approach allowing seamless integration and removal of nodes from the network.

This paper is organized as follows. A description of the Integrated Radio Optical Communications system is provided followed by discussion of Disruption Tolerant Networking. We then address many of the outstanding issue and research topics including addressing, routing, security, scalability and modeling and simulation.

II. Integrated Radio Optical Communications

iROC is a space communications project examining hardware, software and operations level integration. The high level objective is to provide a modern operational optical down-link enhancement to a communications payload for a minimal amount of added size, mass, and power. The mid-level objective is to combine the best features of radio and optical communications into one platform that will unconstrain science missions relative to data capacity. This integration features hybridization in hardware such as where the RF antenna and the optical telescope are combined into a so-called teletenna, which is illustrated in Figure 1.

¹Research Technologist in Data Systems, Architectures, Networks and System Integration Branch (LCA), 21000 Brookpark Road/Mail Stop 54-1

²Aerospace Technologist in Telecommunications, Optics and Photonics Branch (LCP), 21000 Brookpark Road/Mail Stop 54-1, AIAA and Technical Committee Member

^aAlso known as Delay Tolerant Networking. However, in these scenarios, disconnection (disruption) is the key factor resulting

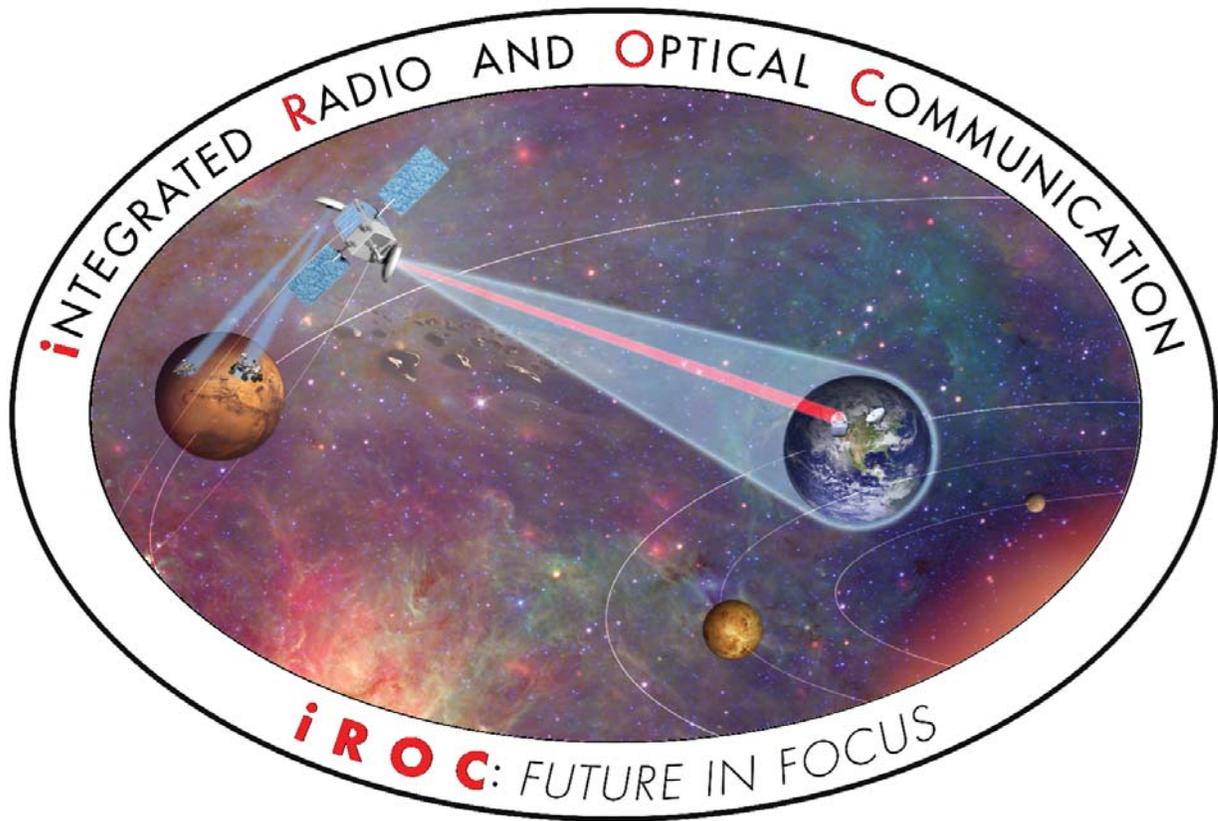


Figure 1. Illustration of the iROc concept operating as a relay in the Mars-Earth system

A beaconless pointing approach is taken in the project; the next evolutionary steps of celestial navigation will enable automated pointing at the precision necessary for deep space optical communications. At deep space distances, the optical beam width guarantees that being pointed optically implies RF pointing. Moreover, the optical beam width is small enough to select specific ground stations (on average, 1550nm from Mars will result in a spot size with the area of Texas) whereas the RF will be ten's of Earth diameters wide.

iROc's two links each feature their own rates, contact times, and channel capacities. iROc is a communications payload for a satellite that offers DTN processing, making it more than a highly capable relay as DTN will make routing decisions: choosing RF or optical is the prime example. Each iROc terminal will have at least one induct, where information flows into the satellite (data, commands, acknowledgments, etc.), and two outducts: the RF and optical transmitters. Each iROc terminal will have its own buffer, contact times, and even security policies.

For iROc deployments, the radio ground stations may not be geographically co-located with their optical complements as the optical ground stations will have additional restrictions based on atmospheric loss. Some of the ground terminals may be in remote locations. Because of the diverse locations of various ground stations and the possible physically different locations of the RF and optical ground stations, difficulties of handovers, load-balancing, data delivery to customers, and scheduling are items that need to be addressed. From the networking standpoint, the most interesting feature of iROc is that it is really two independent down-links.

iROc must be able to meet the data rate needs of the mission, In NASA's Lunar Laser Communications Demonstration (LLCD),² the bus connecting the spacecraft data buffer to the lunar laser space terminal

in the need for store and forward. Propagation delay is simply a mechanism that must be addressed in the hop-by-hop transport protocol

(LLST) operated at 40Mbps, and therefore could not keep up with the LLST when running at any level beyond the minimum.³ For iROC, data rates need to match sensor and science requirement needs. They also need to match the available buses and not over or overwhelm the rest of the system. At the very least, the buffer needs to be able to flow data to the optical radio at its highest rate. To gracefully interface with the rest of the satellite, the bus rates should not impose the bottleneck.

There are many areas of highly asymmetric communications links within the iROC system. For example, at Mars distances, the downlink laser communications system will peak at 267Mbps,⁴ while the Ka-band RF downlink will peak at 84.5Mbps.⁵ Further, uplink communication rates will likely be in the 2kbps range, similarly to the Mars Reconnaissance Orbiter (MRO).⁶

III. Disruption Tolerant Networking

DTN is a generalization of networking as we know it terrestrially. It is a *store and forward* network overlay designed to accommodate disruptions, long propagation delays, and (in theory via late binding^b) mobility. NASA follows RFC4838¹ for the definition of DTN and RFC5050⁷ for the definition of the standard unit of data, the bundle. The bundle consists of a primary bundle block which contains the DTN header, a payload block which holds the payload, and other optional extension blocks that are not defined in the Bundle Protocol specification and reserved for experimental purposes and future use. Protocols such as the Bundle Security Protocol (BSP), RFC6257,⁸ utilize extension blocks.

The *storing* capability is all about holding onto data until it can be forwarded. While there are no established forward-facing links or links with routes to the proper destination, data must be stored and managed. This storage management includes managing storage, checking to see if bundles have expiration and should be removed from the system, and queue management via QoS. For a DTN node, there is finite storage and the possibility that the memory can fill quicker than it is draining and will eventually overflow. QoS assigns priorities to different data allowing for automation of memory allocations. This eases the burden on the equipment and the operator by allowing the DTN forwarding agent to manage the memory. RFC4838¹ assumes memory is inexpensive and plentiful. This may not be the case for deep space systems. Regardless, any buffer has finite size and bandwidth. Clearly too small of a buffer will not accommodate the periods without contact. Too large of a buffer will be prohibitively expensive and complicated. Thus, memory optimization is an area that needs consideration relative to the desired system throughput, the buses available, and link availability.

Forwarding is where the next hop data transfer occurs. The data may be forwarded unreliably with no need for acknowledgment or reliably where some form of acknowledgment is provided by the receiving system.

III.A. Routing

Two types of routing exist in DTN, predictive and opportunistic. Predictive routing is when one knows all the contacts and contact times at any point in time. Predictive routing is also referred to as oracle routing. This requires distribution of contact times to all participating entities. Distribution of the contact graph tables is more a network management function than a routing protocol function. That is, there is no discovery involved. Predictive routing is used in deep space long propagation links where discovery is impractical and orbits and contact times are well known. Opportunistic routing occurs in systems that have non-deterministic contact time. Usually this occurs over links with relatively short propagation delays via some form of discovery protocol.

For deep space deterministic systems the only routing protocol currently available is Contact Graph Routing (CGR). CGR assumes that all nodes in the network are known, all contact times are known, all bit rates are known, and the physical distance between nodes in seconds is known. CGR also assumes that this information is globally shared across the network. This limits CGR's ability to scale.

For non-deterministic systems, and number of routing protocols have been proposed including Delay Tolerant Link State Routers (DTLSR), Probabilistic Routing Protocol using History of Encounters and Transitivity (PRoPHET),⁹ and PRIoritized EPidemic (PREP)¹⁰ to name a few. All of these have some form of discovery. DTLSR attempts to discover the topology of the network to base routing decisions on.

^bLate binding means that the binding of a bundle's destination to a particular set of destination identifiers or addresses does not necessarily happen at the bundle source. Because the destination endpoint identifier (EID) is potentially re-interpreted at each hop, the binding may occur at the source, during transit, or possibly at the destination(s).

PRoPHET was designed for DTNs where there is no fixed topology or schedule. All data forwarding happens at opportunistic encounters between nodes. Patterns in the mobility are used to improve the use of resources as compared with Epidemic Routing to which it is related. PREP prioritizes bundles based on costs to destination, source, and expiry time. Costs are derived from per-link “average availability” information that is disseminated in an epidemic manner. PREP is designed for highly disconnected and mobile networks where no schedule information or repeatable patterns exist.

iROC will be at the nexus between both the predetermined and the ad-hoc network. The laser terminal will have multi-minute long round trip times (RTT)^c. For the Laser and RF link to Earth, a predictive routing mechanisms such as CGR will be used as orbital mechanics, competition, and dollars will determine what ground stations iROC communicates with. The Mars network of satellites and Mars terrestrial assets that iROC services will be close enough to make use of discovery techniques, adding flexibility and scalability. Thus, iROC will need to route between deterministic and non-deterministic (predictive and opportunistic) networks.

III.A.1. Routing Issues

- Presently, there is no addressing in DTN outside of flat naming schemes. This makes route aggregation problematic which makes scaling difficult. Currently, DTN deployments have been in relatively small networks - tens to perhaps hundreds. This scale is indicative of a space-based network. For DTN to be deployed in large networks with millions of nodes, scaling must be addressed.
- Currently, there is no developed capability to redistribute routing information between predictive routing such and CGR and an opportunistic routing protocol.

III.B. Fragmentation

Fragmentation is a useful feature in DTN in order to fully utilize radio links that are intermittent. Two types of fragmentation exist: proactive and reactive. Proactive fragmentation is useful if a contact time and available bandwidth are know. If this is the case, one can mathematically fragment a large bundle into sizes that can be transmitted over the link with that particular link is available. A consequence of the CGR-style information is the provision for proactive fragmentation. If the outgoing queue for the next contact is not full, and it cannot send the next bundle in its entirety, it may fragmented ahead of time to make sure each contact is fully utilized. If a link goes down mid-transmission, the bundle might be reactively fragmented as to minimize retransmission. Reactive fragmentation is used for non-deterministic intermittent links. If one cannot predetermine the contact time and available bandwidth of a particular radio link, one cannot predict the appropriate bundle fragment size. For opportunistic links using the Transmission Control Protocol (TCP) as the transport mechanism, techniques have been devised to determine how much of a bundle has been send during a particular contact and what should be sent during the next opportunity. This has proven to be quite useful.¹¹

III.C. Security

DTN is a network overlay. As such, it has all the security issues associate with any network including authentication, authorization, integrity, confidentiality, and non-repudiation. Furthermore, all underlying networks of which DTN resides above must also be sufficiently protected to enable reliable secure end-to-end communications.¹²

BSP,⁸ outlines security in terms of authentication, integrity, confidentiality and encryption. The BSP realization of these security concepts is still considered experimental and has been shown to induce limitations on the interpretation of DTN. For example, some of the BSP functions add information to the bundle that precedes and follows the payload. Therefore fragmentation is broken - indeed, if a bundle’s security requires pre- and post- payload data, it must be received whole. This is particularly true of the Bundle Authentication Block (BAB), which ensures authenticity and integrity on a single-hop basis. Therefore some research is needed both to determine if the BAB is warranted and if the implementation can be made fragmentation-friendly.

^cRound-trip time is the time required for a signal pulse or packet to travel from a specific source to a specific destination and back again.

There has also been some development into a simplified-BSP.¹³ This is both to tame configuration and network management and to make the code base more manageable. As with BSP, some research is needed to determine if the implementation can be made fragmentation-friendly.

Encryption can be very CPU-intensive, and can be offloaded onto various dedicated hardware platforms. Some research is necessary to determine if encryption is necessary, and if so, if it belongs in the DTN protocol. Clear benefits to hardware implementations include speed, and clear benefits to software implementations include flexibility and customization to achieve the greatest sense of interoperability.

There is a recent firewall-like addition to DTN, named after IPTables. It is called BPTables,¹⁴ where BP stands for Bundle Protocol. It is a rules-based mechanism for determining if a given bundle should be accepted from a certain source or should be forwarded to a certain destination. The rules started off as name-based, but can be easily extended. An example is, “node A cannot accept bundles greater than 5 megabytes.” Some exploration needs to occur to flesh out the implementation - what types of rules are going to be used? What are not? How do they tie in to the rest of the security suite, BSP?

Security Key distribution is a major research area for DTNs. Public Key Infrastructure (PKI)¹⁵ is widely used in terrestrial networks where connectivity and seemingly instant RTTs are more or less guaranteed. PKI suffers when long RTTs are introduced. Consider a key revocation notice sent to a branch of a network that will not receive it for hours, perhaps days. The node that is no longer to be trusted will continue to be trusted until the revocation notice arrives, and so some form of damage control must take place. Obviously terrestrial PKI cannot be used in DTNs.

The outstanding question is “What is the alternative” Are keys required for deep space security? Can observational data be fed into some network management engine (perhaps cognitive) to make real time assertions on confidence? Key distribution, if used, is an area that requires significant research.

Sharing of resources is highly desirable - particularly with regard to space-based system. Consider the international scene at Mars. The Russian Federal Space Agency (RFSA) and the European Space Agency (ESA) already have already been to Mars. The Indian Space Research Organisation (ISRO) has sent an orbiter to Mars. The Chinese National Space Administration (CNSA) has plans to send a rover to Mars in 2020. Each of these agencies has built their own deep space networks throughout the world, and they overlap minimally. Considering the volume of deep space missions and the coverage of the ground, this is sub-optimal; indeed, when ISRO’s Mars Orbiter Mission is behind Mars, the Indian Deep Space Network (IDSN) will be dormant until they launch more missions. Given the ever changing political landscape, it may very well be that CNSA will decide to buy an iROC for their satellites or help NASA complete a NASA project - perhaps CNSA will operate the mission that first flies it. Funding and other unforeseen rocks amidst the flow could push us in this direction. It is therefore worthwhile to consider a complete security concept of operations that enables operations while protecting *us from them* and *them from us*. The best illustration is in the optical domain. Due to atmospheric and weather, there are only so many places that can logically support optical ground stations. At some point mutual distrust and other issues will be outweighed by need for data. Consider when the lives of astronauts depend on it!

III.D. Time and Time Synchronization

Per the RFC5050 specification, bundles are created with a specified lifetime. How this lifetime is determined is not part of the specification and is up to either the application or the deployment. One use of the lifetime value is to purge data from the system once the lifetime value expires. One may implement a deployment where certain data may be considered useful indefinitely and others may have very short lifespans.

DTN works in the framework of absolute time and loose time synchronization. Loose time synchronization implies that systems must be synchronized to some degree. That degree is deployment dependent and largely is relate to the smallest lifetimes of any bundle in the network which is application dependent. For example, if the shortest lifetime bundle lives for 24 hours, than the network probably has to be synchronized to within no less than an hour whereas if the shortest bundle lifetime is 30 seconds, the network probably has to be synchronized to a second or two.

Some assets, such as iROC, will have such advanced atomic clocks and the need for scheduling that time will always be known. Not all communicators in space can be assumed to have such tight timing capabilities, and no Network Time Protocol (NTP) for space has been developed as the problem has been minimally investigated.¹⁶

We live in an age where it is cheaper than ever to launch *things* into space. Indeed, students are even putting Lego robots into low earth orbit, and many universities have CubeSats or other Small Satellite programs. The Time to Live (TTL) for DTN bundles is in seconds from creation, and Internet Protocol (IP) packets are in terms of a hop limit. A similar structure might make more sense in DTN given the wide variations in hardware, and could be implemented in the protocol. Coupled with QoS, it might be sufficient to completely supplant the existing time-based TTL in bundles. The question to answer is: does the limited granularity of several bits of QoS with perhaps a hop limit capture the behavior of time-based TTLs?

III.E. Network Management

Network management is currently an open issue.

Ohio University began initial work on network management in 2009 and release an Internet Draft entitled DING Protocol – A Protocol For Network Management. The Diagnostic Interplanetary Network Gateway protocol (DING). DING is a subscription-based network management protocol designed specifically for use in situations where SNMP’s request / response model does not perform well, e.g. in heavily delayed and / or disrupted networks. This draft has expired with no additional work continuing.

A recent network management service was created for DTN. The proposal was released as an Internet Draft entitled Delay Tolerant Network Management Protocol (DTNMP), which allows for the remote configuration and monitoring of DTN nodes. This is based on the DING work although no reference to DING is provided. The draft has currently expired as of July 2014. Several successful tests were conducted demonstrating that DTNMP *can* be used.

The consensus is that some intelligent control must be realized on remote networks to avoid failure. This promotes the notion of a cognitive engine, or perhaps a distributed (yet delay tolerant) cognitive engine that can use the hooks and data provided from DTNMP to keep the network functional.

IV. Modeling & Simulation

In a space system, the link asymmetries, delays, and overall flux make optimality itself dynamic. To explore advances and to further probe DTN, high-fidelity testbeds are required. In addition, to perform meaningful tests, good traffic model and the connectivity models are required that closely represent a particular deployment scenario.

IV.A. Traffic Model

The traffic model is what stresses the network. Traffic sources can be generated at different scales for different nodes. Both streaming and non-streaming traffic can be implemented depending on the scenario. The QoS requirements and timeliness requirements should also be represented. If one is able to sample existing traffic patterns for a particular or similar deployment, that should be incorporated into the traffic models. Over all, the more closely data can represent reality and applications the better.

A major challenge with DTNs is that there is extremely limited deployment experience. Hence, it is difficult to model DTN traffic.

IV.B. The Connectivity Model

The connectivity model provides the contact times for a given deployment scenario. For deterministic systems such a space-based networks, orbital mechanics makes this relatively easy using tools such as Satellite Orbit Analysis Program (SOAP) or Satellite Took Kit (STK)TM - so long as asset scheduling is not involved. For opportunistic networks, generation of contact times is problematic. One technique is to measure contact times for similar deployment scenarios and incorporate that data. Otherwise, some form of random movement models may be used. As with the traffic model, the better the model represents reality, the better the simulation results.

IV.C. Testing

For deterministic networks, connectivity data can be represented to various fidelities. For example, the bare minimum would only have a list of contact times. A more complete model would be the information that

goes into CGR (bit rates and distances). Other models might include wireless signal rates for RF emulators to run a variety of DTN scenarios against different modulations and encodings. The traffic model would be mission specific and represent all applications one would expect to operate over this network.

For non-deterministic networks one might consider the “Frogger Model,” where a frog (the bundle) hops across logs (links) as they become in reach. It is quite another to consider the case with many frogs of many sizes vying for limited seats on the same logs, and there being many logs going in all directions. The illustrated chaos is actually very repeatable with carefully developed tools.

DTN-for-iROC tests have been conducted and published.²⁰ They were run on a multi-path, multi-hop network with freespace optical links. An STK scenario gave connectivity data that included one-way-light-times. The network behavior was enforced by a network emulator. The traffic model was simple - send various bundles of various sizes and see what arrives. The connectivity model was generated from STK using a program that translated the Azimuth-Elevation-Range reports into CGR and network emulator configurations. The Licklider Transmission Protocol (LTP)²¹ was used. One of the major findings in this experiment was that the particular implementation of LTP was not capable of saturating even slower (10BaseT) links, and therefore while in theory a sound choice, in practice was not useful for modern high-data high-rate communications. LTP uses positive acknowledgments, and the back channel featured enough bandwidth to support the ACKs. Much was learned about the state of DTN at the time. DTN has since grown considerably, and is now ready for the next wave of experiments.²² LTP can also be replaced with similar protocols, such as Saratoga²³ and NACK-Oriented Reliable Multicast (NORM).²⁴

There are several NASA testbeds for DTN. NASA has the DTNBone, which is an open-access DTN network for people to run tests. Recently, there has been a small robotics lab building a collection of Raspberry Pi nodes, robots, and quadcopters. The purpose is for all of these devices to communicate via Bluetooth and to generate realistic ad-hoc connectivity data that can be used to generate the connectivity models we might see at Mars. As each of these nodes can be outfitted with sensors, we can simultaneously create traffic models that will tax iROC’s optical links.

V. Standards

The specification that DTN is based on offer a forum for participation by the general public. They also enforce certain ideals, such as interoperability. Implementations of the same specification must interoperate, and there must be several independently developer implementations. This practice helps weed out poor and ambiguous wording while making sure that other bugs and oversights are carefully inspected. It also means that various organizations implementations will, at least on some established level, work together. Interoperability and openness are necessary to prevent DTN from becoming a singular experiment as opposed to an operating component.

RFC5050 is currently not a standard. It was work that came out of the Internet Research Task Force (IRTF). IRTF does not generate standards; that is the work of the Internet Engineering Task Force. ”RFC5050 defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind.... This RFC is not a candidate for any level of Internet Standard. The IETF disclaims any knowledge of the fitness of this RFC for any purpose and in particular notes that the decision to publish is not based on IETF review for such things as security, congestion control, or inappropriate interaction with deployed protocols.”

There is currently discussion of taking RFC5050 or some other form of Disruption Tolerant Networking protocol(s) to the IETF for standards. A Birds-Of-a-Feather (BOF) meeting will be held in July 2014 at the IETF 90 meeting in Toronto, Ontario, Canada to determine interest and commitment. The results of that meeting should be available by the time of this publication - see www.ietf.org.

DTN as references to RFC5050 is being proposed as a Consultative Committee for Space Data Systems (CCSDS) standard. Work is being performed by the Delay Tolerant Networking Working Group within Space Internetworking Service (SIS) area of CCSDS - see <http://cwe.ccsds.org/sis/>.

VI. Future Research Topics to Consider

The testing of DTN has opened the need for additional research. In addition, there are numerous known problems with the current architecture and protocols that need to be addressed. The following is a list of proposed topics. The purpose is to fuel community-wide discussion and hopefully involve more parties and

philosophies.

VI.A. Known Problems

The following are some of the known problems with the current Bundle Protocol Specification²⁵ (list in no particular order):

- There are no reliability checks required in the DTN bundle. Thus, there is no way to know if a bundle is actually received correctly or even if the primary bundle block is correct. BSP can be used for this, but that is quite a heavy burden to pay - particularly if one wished to deploy DTN in an embedded system as part of a sensor network.
- There is no hop count or equivalent feature to allow one to terminate a bundle if it is stuck in a routing loop.
- Time synchronization is problematic. How one synchronizes 100s, 1000s or millions of nodes in a disjoint and disrupted network is unknown. Also, one cannot assume all systems have accurate clocks.
- Bundle Security Protocol does not work with fragmentation
- Security may break if an intermediary is the security gateway as the bundles may not be transmitted to the destination gateway before arriving at the true destination.
- Expectations from convergence layers is not well defined.
- Expectations of how a DTN bundle agent is suppose to handle items such a no current route but the bundle will not expire for a long time. Should the bundle be discarded held in hopes that a route will appear sometime in the future?
- A flat address space makes scaling and routing difficult
- There is no uniform naming scheme which make routing and discovery problematic
- Mechanisms for identifying, labeling, policing, and otherwise providing QoS for bundle flows are not yet part of the DTN architecture.
- There is no standardized discovery mechanism.
- There are no resolution protocols comparable similar to DNS, ARP, or SIP, and routing protocols with features similar to RIP, OSPF, or BGP have yet been defined. Furthermore, defining mechanisms that are independent of a particular convergence layer adapter or operational environment is a major challenge.

VI.B. Routing & Addressing

Naming and Addressing are essential for developing good routing protocols. Naming and addressing is generally poorly understood. For example within DTN RFC4838, “The concept of using regions as one component of an EID was proposed as a way to distinguish particular subnetworks in an EID. However, this is very problematic in the general sense because it leads to difficulties in multihoming by naming an interface (or set of interfaces) rather than a host (or application process, or set of application processes).”²⁵

Networks are all about aggregating local data into global data, yet this cannot ever be achieved in most DTNs. There have been DTN experiments involving buses and computers distributed to conference attendees to observe the health of a given routing algorithm. It might be more useful to take a step back into the fundamentals and observe connectivity data. As suggested with the above experiments with the robotic testbed and STK/SOAP scenarios, it is possible to create and record connectivity-with-respect-to-time data. These connectivity models can be implemented in Virtual Machines and other network simulators to directly compare routing algorithms against each other rather than doing it in-situ.

Another possibility is to get many time-varying topologies and try to understand persistent patterns in them. If there are topological features that remain more or less constant despite what happens, for example, at the leaves, then this information may beget an addressing scheme. A potential tool is persistent homology,

and other tools of computational topology. It is likely that several addressing schemes would make sense for different styles of DTN (deep space, military, etc.). This should prove an exciting entrance point for mathematics beyond information theory into networking research.

VI.C. Bundle Sizes

There is a multidimensional optimization problem that is largely unexplored: how big should a bundle be? In the iROC case, the up-link will be limited to 2kbps. This up-link has to provide commands, control, and other information in addition to acknowledgments and negative acknowledgments - providing ACKs and NACKs for small segments of data quickly becomes impossible. If a bundle is very small, its payload can be overshadowed by the header. In addition, breaking big data down into tiny bundles means an entry in a routing table for each bundle, which is wasteful of computational resources. Very large bundles solve these issues, but introduce new ones: the recipient might not have the resources to accept large bundles and might be so far away that it cannot negotiate even if a negotiation process existed. This means the recipient can suffer a Denial of Service (DoS) attack. Big bundles may not be completely broadcast before communications is lost. Large bundles also make security more difficult in the event of an upset. Given a contact graph, is enough information present to determine the range of optimal bundle sizes? Of course contact graphs are only applicable to deterministic highly scheduled networks.

VI.D. Network Management Research

Network management of disjoint and disconnected networks is a huge area of research. Besides the protocol development, numerous question need to be consider. The following are just a few examples:

- How does one defines a *metric* for a DTN? Although an overall bit rate of *zero* suggests problems have occurred, one cannot measure the health of a network in bit rate alone. DTN cannot *fix* the physical reality, and cannot be expected to glean all relevant information from the network. Potential measurable health indicators include figuring out how many bundles are transmitted versus how many are discarded, and if bundles are discarded based on TTL (hop count or otherwise).
- Is end-to-end bundle delivery time meaningful? After all, if the bundle lifetime is set to one day and the destination receives the bundle prior to it's expiration time, is that not that a successful transmission? If it is sent an hour later than it could have been, this could be a sign of system health if it was preempted by a higher priority bundle.⁴
- Can one add a cognitive engine^d into the system in order to learn what information is truly important to report back and what information can be inferred from other information?

VI.E. Network Coding

There may be applications of network coding to entire DTNs or particular portions of them and much research is ongoing^{26,27} In addition, an Internet Draft (currently expired) entitled "Random Binary FEC Scheme for Bundle Protocol" has been submitted on the subject. When assumptions regarding connectivity, availability, and RTTs must be relaxed, the type (linear/nonlinear) and breadth of application of network coding may change considerably. This analysis may lead to a rigorous framework for disconnected networks.

VII. Final Thoughts

The individual components of DTN have made great strides in recent years. There is still fertile ground for fundamentals research. There are also many research projects ahead at various levels of integration - both for internal systems and systems of systems.

What does the future hold for iROC? The culmination of the above takes the form of multiple iROCs. Instead of a single-homed network, where iROC is a privileged node that connects the Martian and Earth

^dA cognitive system will make mistakes - mistakes are part of learning. Given the space industries predisposition towards safety and outright determination, how could such a method of operation be adopted? In the event of failure, how (and should) blame/fault be determined?

networks, there is a multi-homed network. There may also be iROC swarms around multiple planets supporting disjoint missions - the common intersection in this case is the ground segment. This large type of network is a very long term goal, and really stresses the flexibility of networking software and the need for interoperable standards. What is really being tested here is a multi-decade vision where nodes and architectures get added to the overall communications system as time goes on while respecting existing nodes. Consider how the deep space network still services the Voyager I and II spacecraft - they continue to send messages to the Deep Space Network on Earth to receive routine commands and return data. The Voyager mission contact schedules must be taken into account when newer assets attempt to use the Deep Space Network. The point is that certain missions may greatly outlive others and in some form they must all work in concert or not at all. Some planning needs to happen to at least form a baseline for how we make the leap from point-to-point communications to a network.

References

- ¹Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and Weiss, H., "Delay-Tolerant Networking Architecture," RFC 4838, Internet Engineering Task Force, April 2007.
- ²NASA, "Lunar Laser Communications Demonstration," <http://esc.gsfc.nasa.gov/267/271.html>, 2014.
- ³Boroson, D. M., Robinson, B. S., Murphy, D. V., Burianek, D. A., Khatri, F., Kovalik, J. M., Sodnik, Z., and Cornwell, D. M., "Overview and results of the Lunar Laser Communication Demonstration," 2014.
- ⁴Daniel E. Raible, Robert R. Romanofsky, J. M. B. J. M. N. A. G. H. A. J. S. A. L. N., "On the Physical Realizability of Hybrid RF and Optical Communications Platforms for Deep Space Applications," 2014.
- ⁵Daniel J. Zeleznikar, Jennifer M. Nappier, J. A. D., "Ka-band Link Study and Analysis for a Mars Hybrid RF/Optical Software Defined Radio," 2014.
- ⁶Fong, M., "Mars Reconnaissance Orbiter Cost Analysis Documentation of Requirements (CADRe) (Launch Submittal)," 2011.
- ⁷Scott, K. and Burleigh, S., "Bundle Protocol Specification," RFC 5050, Internet Engineering Task Force, Nov. 2007.
- ⁸Symington, S., Farrell, S., Weiss, H., and Lovell, P., "Bundle Security Protocol Specification," RFC 6257, Internet Engineering Task Force, May 2011.
- ⁹Lindgren, A., Doria, A., and Schelén, O., "Probabilistic routing in intermittently connected networks," *ACM SIGMOBILE mobile computing and communications review*, Vol. 7, No. 3, 2003, pp. 19–20.
- ¹⁰Ramanathan, R., Hansen, R., Basu, P., Rosales-Hain, R., and Krishnan, R., "Prioritized epidemic routing for opportunistic networks," *Proceedings of the 1st international MobiSys workshop on Mobile opportunistic networking*, ACM, 2007, pp. 62–66.
- ¹¹Ivancic, W., Paulsen, P., Stewart, D., Eddy, W., McKim, J., Taylor, J., Lynch, S., Heberle, J., Northam, J., Jackson, C., et al., "Large File Transfers from Space using Multiple Ground Terminals and Delay-Tolerant Networking," *Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE*, IEEE, 2010, pp. 1–6.
- ¹²Ivancic, W. D., "Security analysis of DTN architecture and bundle protocol specification for space-based networks," *Aerospace Conference, 2010 IEEE*, IEEE, 2010, pp. 1–12.
- ¹³Birrane, E., "Streamlined Bundle Security Protocol Specification," Internet-Draft draft-irtf-dtnrg-sbsp-01, Internet Engineering Task Force, 05 2027, Work in progress.
- ¹⁴Iannicca, D. and Hylton, A., "BPTables Firewall Concept of Operations," Tech. rep., To Be Published September 2014.
- ¹⁵Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and Polk, W., "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," RFC 5280, Internet Engineering Task Force, May 2008.
- ¹⁶Mills, D. and Nair, H., "Timekeeping in the Interplanetary Internet," 2004.
- ¹⁷J Burgess, B Gallagher, D. J. B. L., "MaxProp: Routing for vehicle-based disruption-tolerant networks," 2006.
- ¹⁸Aruna Balasubramanian, B. N. L. and Venkataramani, A., "DTN Routing as a Resource Allocation Problem," 2007.
- ¹⁹Pan Hui, Augustin Chaintreau, J. S. R. G. J. C. C. D., "Pocket Switched Networks and Human Mobility in Conference Environments," 2005.
- ²⁰Hylton, A. G., Raible, D. E., Juergens, J., and Iannicca, D., "On Applications of Disruption Tolerant Networking to Optical Networking in Space," 2012.
- ²¹Kruse, H., Jero, S., and Ostermann, S., "Datagram Convergence Layers for the Delay- and Disruption-Tolerant Networking (DTN) Bundle Protocol and Licklider Transmission Protocol (LTP)," RFC 7122, Internet Engineering Task Force, March 2014.
- ²²Pöttner, W.-B., Morgenroth, J., Schildt, S., and Wolf, L., "Performance comparison of DTN bundle protocol implementations," *Proceedings of the 6th ACM workshop on Challenged networks*, ACM, 2011, pp. 61–64.
- ²³Wood, L., Eddy, W., Smith, C., Ivancic, W., and Jackson, C., "Saratoga: A Scalable Data Transfer Protocol," Internet-Draft draft-wood-tsvwg-saratoga-15, Internet Engineering Task Force, 04 2019, Work in progress.
- ²⁴Adamson, B., Bormann, C., Handley, M., and Macker, J., "Negative-acknowledgment (NACK)-Oriented Reliable Multicast (NORM) Protocol," RFC 3940, Internet Engineering Task Force, Nov. 2004.
- ²⁵Wood, L., Eddy, W. M., and Holliday, P., "A bundle of problems," *Aerospace conference, 2009 IEEE*, IEEE, 2009, pp. 1–17.
- ²⁶Widmer, J. and Le Boudec, J.-Y., "Network coding for efficient communication in extreme networks," *Proceedings of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*, ACM, 2005, pp. 284–291.

²⁷Balasubramanian, A., Levine, B., and Venkataramani, A., “DTN routing as a resource allocation problem,” *ACM SIGCOMM Computer Communication Review*, Vol. 37, No. 4, 2007, pp. 373–384.