Contact Graph Routing in DTN Space Networks: Overview, Enhancements and Performance

Giuseppe Araniti, Nikolaos Bezirgiannidis, Edward Birrane, Igor Bisio, Scott Burleigh, Carlo Caini, Marius Feldmann, Mario Marchese, John Segui, and Kiyohisa Suzuki

Giuseppe Araniti is with the University Mediterranea of Reggio Calabria.

Nikolaos Bezirgiannidis is with the Democritus University of Thrace.

Edward Birrane is with Johns Hopkins University Applied Physics Laboratory.

Igor Bisio and Mario Marchese are with the University of Genoa.

Scott Burleigh is with the California Institute of Technology.

John Segui is with Speedy Packets, Inc. and Pasadena City College.

Carlo Caini is with the University of Bologna.

Marius Feldmann is with Technische Universität Dresden.

Kiyohisa Suzuki is with Japan Aerospace Exploration Agency.

ABSTRACT

Delay- and Disruption Tolerant Networks (DTNs) are based on an overlay protocol and on the store-carry-forward paradigm. In practice, each DTN node can store information for a long time before forwarding it. DTNs are particularly suited to cope with the challenges imposed by the space environment. This paper is focused on routing in space DTNs, and in particular on contact graph routing (CGR) and its most representative enhancements, available in the literature, which are briefly surveyed in this work.

Moreover, the applicability and the obtained performance of the DTN protocol stack and of the CGR have been evaluated by presenting results from real experimental experiences such as the Deep Impact Network experiment (employing the EPOXI space cruise), the JAXA jointly performed space link demonstrations with NASA (where the JAXA's GEO relay satellite called Data Relay Test Satellite has been used), the Space Data Routers European Project, and the pilot operation of a DTN implementation on the International Space Station (ISS).

INTRODUCTION TO SPACE NETWORKS AND DTNS

Every mission into deep space has a communications system to carry commands and other information from Earth to a spacecraft or to a remote planet and to return scientific data to Earth [1]. Communications systems are central to the success of space missions. Large amounts of data need to be transferred (for example, nearly 25 TB in 2013 concerning the Mars Reconnaissance Orbiter (MRO)), and the demand will grow in the future [1] because of the employment of more sophisticated instruments that will generate more data. This will require the availability of high network transfer rates. Satellite systems already have to cope with difficult communication challenges: long round trip times (RTTs); the likelihood of data loss due to errors on the communication link; possible channel disruptions; and coverage issues at high latitudes and in challenging terrain. These problems are magnified in space communications characterized by huge distances among network nodes, which imply extremely long delays and intermittent connectivity. At the same time, a space communications system must be reliable over time due to the long duration of space missions. Moreover, the importance of enabling Internet-like communications with space vehicles is increasing, realizing the concept of extended Future Internet, an IP (Internet Protocol) pervasive network of networks including interplanetary communication [2], where a wide variety of science information values are acquired through sensors and transmitted.

The Delay- and Disruption Tolerant Network (DTN) architecture [3] introduces an overlay protocol that interfaces with either the transport layer or lower layers. Each node of the DTN architecture can store information for a long time before forwarding it. Thanks to these features, a DTN is particularly suited to cope with the challenges imposed by space communication. As summarized in [4], the origin of the DTN concept lies in a generalization of requirements identified for interplanetary networking (IPN), where latencies that may reach the order of tens of minutes, as well as limited and highly asymmetric bandwidth, must be faced.

However, other scenarios in planetary networking, called "challenged networks," such as military tactical networking, sparse sensor networks, and networking in developing or otherwise communications-challenged regions, can also benefit from the DTN solution. Delays and disruptions can be handled at each DTN hop in a path between a sender and a destination. Nodes on the path can provide the storage necessary for data in transit before forwarding it to the next node on the path. In consequence, the contemporaneous end-to-end connectivity that Transmission Control Protocol (TCP) and other standard Internet transport protocols require in order to reliably transfer application data is not required.

In practice, in standard TCP/IP networks,

 0163-6804/15/$25.00 © 2015 IEEE

which assume continuous connectivity and short delays, routers perform non-persistent (short-term) storage and information is persistently stored only at end nodes. In DTN networks, information is persistently (long-term) stored at intermediate DTN nodes. This makes DTNs much more robust against disruptions, disconnections, and node failures.

The Bundle Protocol (BP) [5] is a key element of the DTN architecture, where the basic unit to transfer data is a bundle, a message that carries application layer protocol data units, sender and destination names, and any additional data required for end-to-end delivery. The BP can interface with different lower layer protocols through convergence layer adapters (CLAs). CLAs for TCP, UDP, Licklider Transmission Protocol (LTP), Bluetooth, and raw Ethernet have been defined. Each DTN node can use the best suited CLA for the forwarding operation.

BP provides useful features such as:

- Custody transfer, where an intermediate node can take custody (i.e. responsibility) of a bundle, relieving the original sender of the bundle which might never have the opportunity to retransmit the application data due to limited resources.
- Proactive and reactive bundle fragmentation, the former to tackle intermittent periodic connectivity when the amount of data that can be transferred is known a priori, the latter, which works a posteriori, when disruptions interrupt an ongoing bundle transfer.
- Late binding, where, for example, when a bundle destination endpoint's identifier includes a dynamic name server (DNS) name, only the CLA for the final DTN hop might have to resolve that DNS name to an IP address, while routing for earlier hops can be purely name based.

Routing is a critical problem in DTN networks. Quoting from [6], "*the routing objective of traditional routing schemes has been to select a path which minimizes some simple metric (e.g. the number of hops). For DTN networks, however, the most desirable objective is not immediately obvious.*" Nodes are not constantly connected. Storage and energy management affect DTN routing. A possible aim may be increasing the probability of bundle delivery, but also reducing the delivery delay may be important. Routing over DTN networks deserves close attention and is the object of the next section. A suitable solution for space networks is represented by Contact Graph Routing (CGR) [7], where each node on the path computes a route from itself to the bundle destination based on a computed graph. We include a brief tutorial on basic CGR, highlight CGR issues and enhancements, and summarize CGR performance over space networks. Finally, the conclusions are drawn.

Routing in DTN Space Networks

Routing and Forwarding in DTN vs the Internet

Given the aforementioned challenges of space communication, it is not surprising that the methods used for computing routes in a space network should be different from those used in Internet routing. To aid in explaining these differences, it may be helpful to return briefly to first principles.

In general, we might say that *routing* is the procedure by which we select the best path for conveying data from source node A to destination node Q in a network. Routing would be trivial if every node could simply transmit directly to every other, but for large networks this is not possible. In recognition of this complexity, a network host plans a route for a data item before issuing it. The network state information on which this planning is based includes the network's "topology" and a list of all known connections between nodes. In a DTN-based network, this list may include additional information such as the speed of each connection and perhaps the storage capacity of each node.

However, network state information may change over time while traffic is traversing the network, and therefore the most efficient route may change while data is en route. For this reason, routing may occur at every branch point to take advantage of newly available information, and consequently it is more accurate to say that *routing* is the procedure by which, at each point in the path from A to Q, we select a neighboring branch point to transmit the data to, believing that branch point to be on the best path for conveying the data to its destination. To make this selection, we may compute a new route based on the network state information currently available at this point or we may simply continue along the path previously computed by another node.

In the Internet this selection can be done with high confidence because information about changes in network state information can be propagated so quickly that each node's current understanding of the state of the network is almost always correct. That understanding may be incomplete, because routing in the network may be compartmentalized: the network state information exposed to any node may be limited to nodes in the local "domain" (including nodes that are on the border between the local domain and adjacent domains that serve as "gateways" between domains). Nonetheless, routing decisions can be made confidently in the expectation that the distribution of network state information within other domains is as rapid and comprehensive as within the local domain. Each node is continuously connected to a small number of neighboring nodes; routing is simply a matter of choosing the neighboring node that's on what seems to be the best path.

In a space network, or in one of the previously mentioned "challenged" networks where DTNs are applied, this is not true: since connectivity is intermittent and/or signal propagation times are long, changes in the network state may occur more rapidly than information about those changes can be propagated. Routing is still a matter of choosing a neighboring node to transmit directly to, but determination of the best path is constrained by lack of knowledge of the current state of the network, and it may not be possible to transmit immediately to the neighboring node that is the nearest branch point on the best path.

Given the aforementioned challenges of space communication, it is not surprising that the methods used for computing routes in a space network should be different from those used in Internet routing. To aid in explaining these differences, it may be helpful to return briefly to first principles.

CGR is a dynamic algorithm that computes routes through a time-varying topology of scheduled communication contacts in a DTN network. It can be successfully applied not only to an Interplanetary Internet, but also to LEO satellite communications, as in both cases link availability is known a priori.

SURVEY OF CURRENT WORK

Strategies for dealing with these obstacles have been the focus of most DTN research for longer than a decade. A key discriminator among these strategies is the assumed timeliness and accuracy of the network state information available to every node in the network. Several surveys of DTN routing schemes have been conducted, and a hierarchy of DTN routing approaches, ranging from those with zero configuration information to those with perfect knowledge of the network, has been defined. Approaches that assume minimal accurate network state information have historically been considered "opportunistic" while those that assume complete network state information are regarded as "deterministic."

Significant algorithms belonging to the category of opportunistic approaches include single-hop multi-cast forwarding (Spray and Wait), in-network exchange of link information (DTLSR), and probabilistic analysis of predicted node contact (PRoPHET). All of these rely on the exchange of infrastructure and/or in-network measurements in a timely manner to support on-demand calculations of routes and forwarding hops. Opportunistic approaches often apply a replication-based (alternatively, "flooding-based") strategy. Using this strategy, messages are typically duplicated either a fixed number of times or else a variable number of times based on contact probability. In networks with high node mobility and nearly random contact establishment, the delivery success rate of this class of approaches is higher than approaches that rely on the accuracy of current network state information.

On the other hand, in networks where contacts are predictable, the more deterministic algorithms can achieve high rates of delivery success with less waste of bandwidth and buffer space. Algorithms such as MARVIN and Contact Graph Routing [7] belong to this second category. Accurate contact predictions are distributed to the nodes in the network, enabling network graphs to be built and used to make routing decisions on a hop-by-hop basis. MARVIN encodes information about the operational environment (planetary ephemeris data) and infers contact opportunities from this knowledge. Similarly, the numerous MANET routing approaches also base their operation on evolving graphs. The Contact Graph Routing (CGR) algorithm is a formulation of the perfect knowledge approach, and is currently being extended to work in less-perfect knowledge systems. CGR is discussed in more detail in the next section.

CONTACT GRAPH ROUTING (CGR)

CGR is a dynamic algorithm that computes routes through a time-varying topology of scheduled communication contacts in a DTN network. It can be successfully applied not only to an Interplanetary Internet, but also to LEO satellite communications, as in both cases link availability is known a priori. However, this perfect knowledge does not reduce the complexity of the route computations, as CGR must consider that links among nodes in the network change over time. For an exhaustive explanation we refer the reader to the CGR section of the ION Design Guide (the Interplanetary Overlay Network (ION) implementation of DTN, including the Design Guide, is available at https://sourceforge.net/projects/ion-dtn/) or to the CGR Internet Draft [7]. Here we provide only a few key points of CGR's functionality.

The basic strategy of CGR is to take advantage of the fact that, since space flight communication operations are planned in detail by mission operators, the communication routes between any pair of "bundle agents" in a population of nodes, all of which have been informed of one another's plans, can be inferred from those plans rather than discovered via dialogue.

The foundation of contact graph routing is the "contact plan," a time-ordered list of scheduled, anticipated changes in the topology of the network. The entries in this list are termed "contacts"; each one is an assertion that transmission from node X to node Y at nominal data rate R will begin at time T1 and will end at time T2. Note that this assertion implicitly also defines the "volume" (or "capacity") of the contact, which is the maximum amount of data that can be transferred during the contact, given by the product of contact length (T2 – T1) and nominal transmission rate R.

Each node uses the contacts in the contact plan to build a "routing table" data structure. A routing table is a list of "route lists," one route list for every possible destination node in the network. Each route in the route list for node D identifies a path to destination node D, from the local node, that begins with transmission to one of the local node's neighbors in the network, the initial receiving node for the route, termed the route's "entry node." The route list entry for each neighbor contains the best route that begins with transmission to that neighbor. Also noted for each route are:

- All of the other contacts that constitute the remaining segments of the route's end-to-end path.
- The estimated "cost" of this route, e.g. the end-to-end delivery latency.
- The "forfeit time" for this route, i.e. the latest time by which the bundle must have been forwarded to the route's entry node in order to have any chance of traversing this route.

To compute a new route list for node D:

- We construct an abstract contact graph, a directed acyclic graph whose root is a notional contact from the local node to itself and whose other vertices are all other contacts that can contribute to some end-to-end path to D. A terminal vertex is also included in the graph, constituting a notional contact from node D to itself.
- We perform a series of Dijkstra searches within this graph. On each search we find the lowest-cost route that begins at the root of the graph and ends at the terminal vertex. Each time a route is computed, we add it to the node's list of routes and then remove that route's initial contact from the contact graph before searching for the next best route. The search series is terminated as soon as a search fails to find a route.

Note that the routes in the route list need not be continuous. Each segment of the path is an opportunity to send data from node X to node Y; once a bundle has reached node Y it may well reside in storage at node Y for some length of time, awaiting the start of the opportunity to be forwarded from node Y to node Z, and so on.

So when a bundle must be transmitted from node A to node Q we consult the route list for node Q. Some of the routes in the list may be unusable. For example, a route may be temporarily unavailable (transmission to the entry node is "blocked" due to a detected or asserted loss of connectivity); or the best-case delivery time on a route may be greater than the bundle's time-to-live (the bundle would be purged before delivery); or the "residual capacity" of the initial contact on the route (the capacity that has not been allocated yet to higher-priority bundles) may not be enough to contain the bundle. Note that this latter check is a form of embryonic congestion control: a route is considered unusable if its first contact is already fully subscribed, causing the bundle to be redirected to less congested routes.

Of the usable routes, we choose the one with the lowest cost and queue the bundle for transmission to that route's entry node. If the list of bundles queued for transmission on some route is non-empty at the time that route's forfeit time is reached, new routes must be computed for all of those bundles.

The key advantage of CGR is that, like Internet routing, it can be done with high confidence, as it is based on accurate information about the network's topology. The difference is that:

- The topology on which routing is based is not the currently known current topology but rather an anticipated time-varying topology.
- Since changes in the network's topology are scheduled in the course of mission planning, information about those changes can be propagated long before they occur. Just as in the Internet, each node's understanding of the topology of the network at any moment is almost always correct: while propagation of information about network topology changes is slow, it is still "faster" than the rate at which the changes themselves occur.

So again, routing is a matter of choosing a neighboring node to transmit directly to. Again, it may currently be impossible to transmit to the neighboring node that is the nearest branch point on the best path, but at least determination of the best path is possible because topology knowledge is generally accurate.

CGR Issues and Enhancements

Ever since CGR first appeared, the research community has worked on improving its functionality and usage. For instance, path selection with Dijkstra's algorithm, proposed as "Enhanced CGR," has now become part of the core CGR functionality. Research activity on CGR is still very active and further enhancements have been proposed to cope with residual issues. A short list of the most representative is presented below, divided into short-term modifications to the algorithm and long-term prospects for CGR evolution.

Short Term Evolution

In route computation classical CGR assumes that bundles will be sent at the contact start time or, if the contact is currently in progress, immediately. That is, it does not consider the queueing delay caused by other bundles in the outbound buffer waiting for transmission. For this reason, a modified version of the CGR algorithm, namely CGR-ETO, was introduced in [8] to incorporate the available queue length information. CGR-ETO utilizes the earliest transmission opportunity (ETO) contact parameter, the earliest plausible time that a bundle of a specific priority can be forwarded during this contact, replacing contact start time with ETO during contact graph traversals. Queue length information can be easily obtained at the local node and is updated upon bundle routing. Obtaining useful (i.e. not obsolete) queue length information from other nodes is challenging and requires the transmission of update messages, e.g. using the Contact Plan Update Protocol (CPUP) [8].

A bundle may be assigned to a route that is already fully subscribed, provided that the bundle's priority is higher than that of some of the bundles currently assigned to that route. For this reason CGR does not take into account bundles of lower priority in the "residual volume" computation check. The contact oversubscription that derives from this policy is informally called contact "overbooking." The aim of the overbooking management adaptation is to mitigate as much as possible the consequences of this contact oversubscription.

In an overbooking example of a future contact, some low priority bundles put in the queue to proximate node X will miss their contact, to accommodate higher priority bundles. This situation is tackled by standard CGR a posteriori, by re-forwarding the "bumped" bundles once their forfeit time expires (usually at the overbooked contact's end-time). This handling, although robust, is not efficient. By contrast, overbooking management acts a priori, by re-forwarding as soon as possible any bundles that are destined to miss the contact, i.e. immediately after forwarding the higher priority bundle that has caused the oversubscription. Results presented in [9] show that overbooking management and CGR-ETO are complementary and effective in improving routing decisions.

Long Term Evolution

Path Encoding CGR Extension — The standard CGR model computes a feasible path through the network and uses that to select the most appropriate next step in the routing process. The Path Encoding CGR extension takes that calculated path and attaches it to the message. Downstream nodes may then merely verify the continued feasibility of the encoded path rather than calculate a new path from scratch at every hop in the network [10]. This approach yields four benefits. First, paths are "re-used" as long as they are verified against local knowledge at downstream nodes, thereby avoiding a complex route calculation at every hop in the network. This is a particularly important optimization when implementing routing decisions on resource-

So again routing is a matter of choosing a neighboring node to transmit directly to. Again it may currently be impossible to transmit to the neighboring node that is the nearest branch point on the best path, but at least determination of the best path is possible because topology knowledge is generally accurate.

Due to its flexibility, CGR can be enhanced in order to be applied not only to deterministic but also to opportunistic scenarios. This would allow its application as core routing in large-scale DTN deployments with various, heterogeneous contact types.

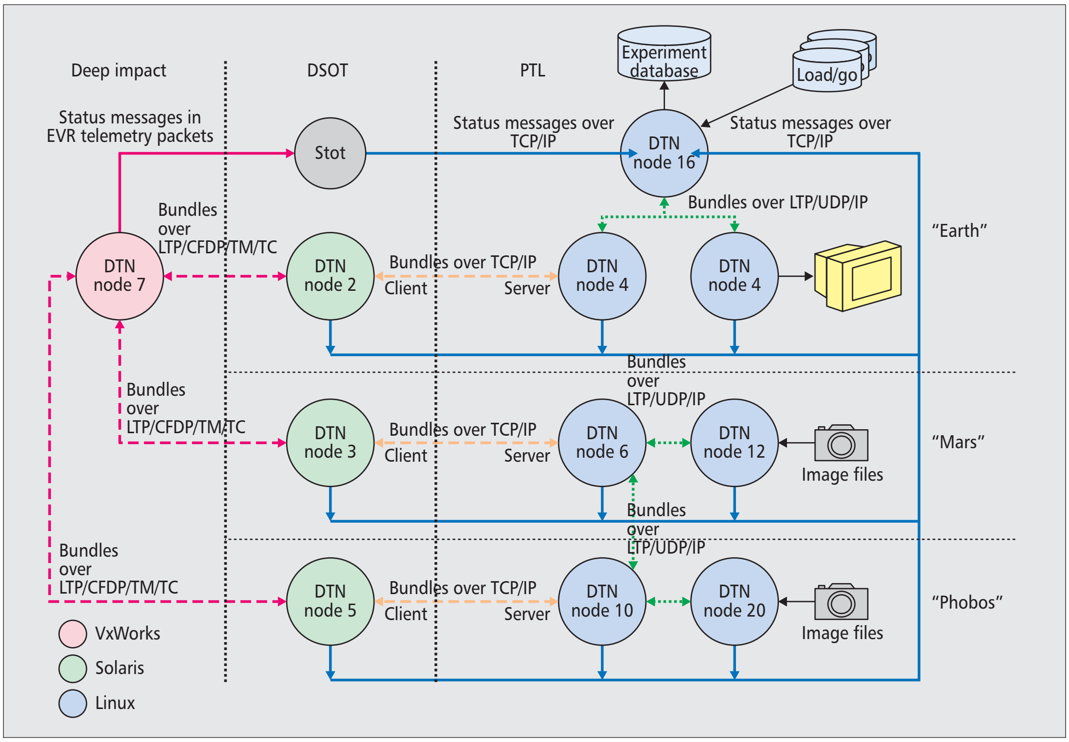

Figure 1. DINET network topology.

constrained flight processors. Second, an encoded path simply needs to remain feasible to be validated, even if a potentially better path could be recalculated. Honoring pre-computed but potentially suboptimal paths provides a natural damping function that resists routing loops in networks undergoing topological changes or congestion. Third, supporting feasible-versus-optimal allows the use of novel cost functions for route selection algorithms that can optimize network utilization rather than individual message delivery cost, which is an important balance in space-based sensing constellations. Fourth, sending path information with a message provides meta-data exploitable for new research such as path-based congestion prediction and topological synchronization.

Opportunistic CGR Extension — Due to its flexibility, CGR can be enhanced in order to be applied not only to deterministic scenarios but also to opportunistic scenarios. This would allow its application as core routing in large-scale DTN deployments with various, heterogeneous contact types, such as scenarios including both terrestrial and space nodes, thus leading to a unified DTN routing approach.

To this end, a "divide and conquer" strategy can be envisaged, where the overarching routing mechanism follows a hybrid (deterministic/opportunistic) approach. The overall network is decomposed into different *regions* where either standard CGR or a modified version able to forward bundles in a probabilistic way is used, depending on the link characteristics (deterministic/opportunistic) of the region. By contrast, regions are always interconnected by standard CGR, as their connectivity is reasonably assumed deterministic.

A CGR evolution proposal to tackle opportunistic forwarding consists of the following steps: first, the contact plan is extended to include contacts with a probability of occurrence lower than 1; second, routes to a destination are calculated as before, thus ignoring the contact probabilities; finally, copies of a message are forwarded to the entry nodes of all opportunistically discovered routes that increase the message's aggregate delivery probability by more than a given threshold. The algorithm is designed to throttle back the number of copies automatically as the aggregate expectation of delivery success on the selected routes increases.

After this short overview of ongoing research, we will address selected experiments that made use of CGR.

Experiments over Space Networks

This section briefly describes four experiences in real space networks carried out by space agencies or within the framework of international projects. These experiences aim at investigating, in general, the effectiveness of the DTN paradigm over operational space networks and, in particular, the effectiveness of the CGR algorithm.

During the DINET experience, NASA performed a first attempt to test their DTN implementation over a real system. The transmissions of photos from remote planets, detailed in the following paragraph, were successfully completed and CGR performed quite satisfactorily. The DRTS DTN project, recently completed by NASA and JAXA, confirmed that it is feasible to use DTN with CGR in real spacecraft operations. The Space-Data Routers project, a European Commission funded initiative, demonstrated that CGR can contribute to efficient data dissemination. Finally, for the sake of completeness, NASA's experience with the ION implementation on the International Space Station (ISS) is described. While this topology was too simple to evaluate the effectiveness of CGR in a large network, the obtained results highlighted important benefits obtained by the application of the DTN paradigm, as listed in detail below.

The Deep Impact Network Experiment

The Deep Impact Network (DINET) project was an experimental validation of "ION" (Interplanetary Overlay Network), JPL's implementation of the DTN protocols. The ION software, including the first implementation of Contact Graph Routing, was uploaded to the backup flight computer of the EPOXI (formerly Deep Impact) spacecraft on 18 October 2008, and was operated continuously from that date until 13 November 2008.

EPOXI was at that time in an inactive cruise period while en route to encounter comet Hartley 2 (in November 2010). The one-way signal propagation time from EPOXI to Earth was initially 81 seconds, dropping to 49 seconds by the end of the four-week exercise. The spacecraft was between 9.1 million and 15.1 million miles from Earth during the experiment.

Uploading the ION software to EPOXI enabled the spacecraft to function as a DTN router in an 11-node network (Fig. 1).

The spacecraft was assigned node number 7 for this exercise and was the only node of the network that was not physically resident at the Jet Propulsion Laboratory (Pasadena, California). Nodes 2, 4, 8, and 16 played the role of "Earth" in the experiment; nodes 3, 6, and 12 functioned as a notional "Mars"; nodes 5, 10, and 20 impersonated "Phobos." The Mars and Phobos nodes simulated the acquisition of images and the transmission of those images back to Earth via the EPOXI spacecraft, which acted as a relay router in space. Each dashed line in the topology diagram represents a sequence of DTN network contacts; the solid blue lines indicate the out-of-band local area network used to instrument the experiment. Note that the topology of the DINET network included cross-links between nodes 6 and 10 (Mars and Phobos), providing alternative paths for data to and from nodes 12 and 20; that is, the DINET network was not a simple tree, so the CGR route selection decisions made at nodes 6 and 10 were non-trivial.

Over the course of the four weeks of flight testing, the DTN software reliably conveyed 292 images (about 14.5 MB) through the network, together with command traffic from the Earth nodes to the Mars and Phobos nodes. No data were lost or corrupted anywhere in the network, and ground station handovers and transient failures in the Deep Space Network uplink service were handled automatically and invisibly. CGR generally performed well, but several bugs in the initial ION implementation resulted in some under-utilization of network capacity. Those bugs were addressed in later versions of ION.

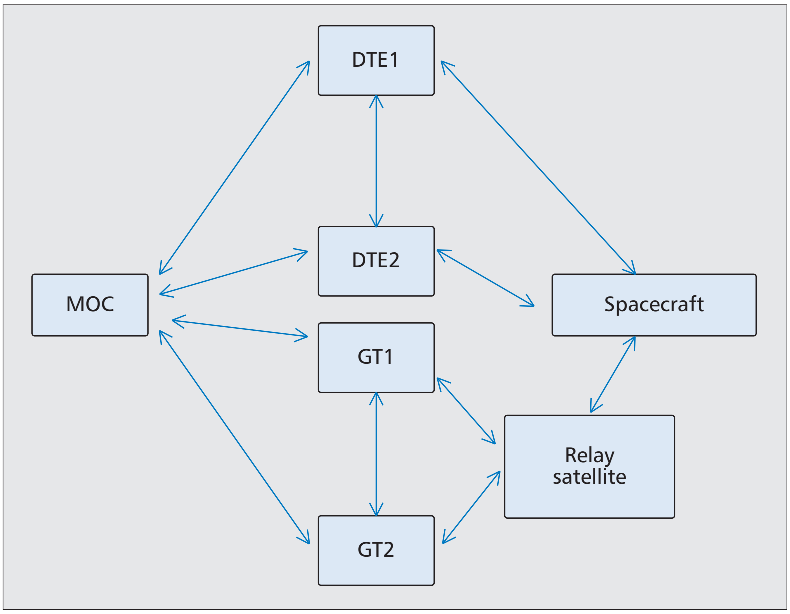

Figure 2. Topology of the JAXA DRTS space experiments.

JAXA DRTS Testing

For the purpose of studying the feasibility of autonomous routing and high-integrity data forwarding in existing and future anticipated space network architectures, JAXA jointly performed a series of experimental tests with NASA in 2012-2013 to evaluate the DTN architecture and CGR [11].

JAXA's GEO relay satellite "Data Relay Test Satellite (DRTS)" and its tracking stations were used in this measurement campaign. The data relay space link is referred to as the "inter-satellite link" in the following discussion. In the tests, ION was used on all nodes to evaluate the performance of BP, LTP, and CGR. Several network topologies were investigated, including direct connectivity between a LEO spacecraft and a ground network and relayed communications between a remote planetary surface and earth's surface connected via a relay spacecraft.

The topology considered here, by contrast, is typical of an earth observation mission. It consists of the following seven DTN nodes shown in Fig. 2: one mission operation center (MOC); one LEO spacecraft; one GEO satellite acting as a relay; two ground terminals in between the GEO satellite and the MOC (GT1 and GT2); and two direct to earth ground stations (DTE1 and DTE2) in between the LEO and the MOC. Data generated on board the LEO satellite can reach the MOC either passing through the GEO relay (and then either via GT1 or GT2) or via DTE1 or DTE2. The task of CGR is to dynamically find the best route, taking into account link intermittency.

The test conditions included the actual signal propagation and processing delay over the DRTS's inter-satellite link and link intermittency, either scheduled (e.g. due to orbital mechanics) or random (e.g. due to space link failures). Contact plans were developed based on the actual resource allocation plan for the inter-satellite link. The obtained results are presented

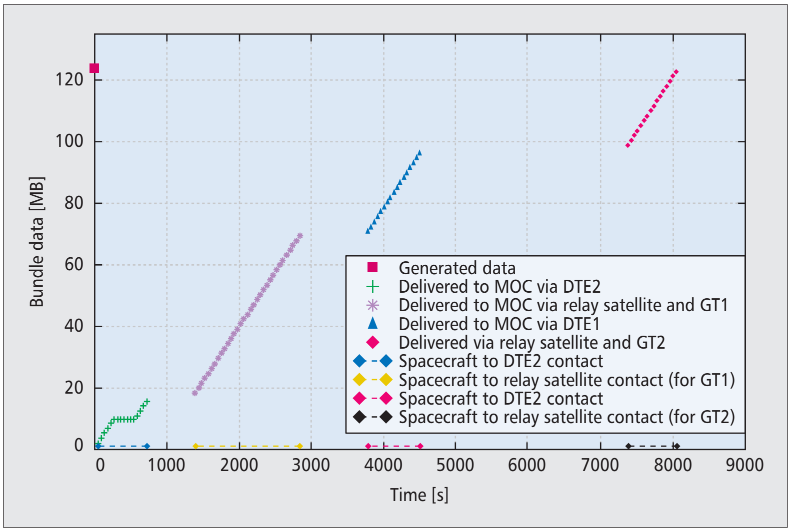

Figure 3. An example of data transfer from a LEO spacecraft and the MOC, via alternative routes dynamically selected by CGR on the basis of the contact plan. All data are routed in accord with contact information, as expected.

in Fig. 3, which shows that all data were automatically transferred to the MOC via four routes, selected by CGR conforming to the contact plan, as expected. These results attest to the feasibility of using DTN with CGR for autonomous routing and reliable data forwarding in real spacecraft operations.

Space Data Routers

The applicability of CGR in the ground segments of space missions was evaluated in "Space-Data Routers" (www.spacedatarouters.eu), a European FP7 project that exploited the DTN architecture to improve the dissemination of space mission data with respect to volume, timeliness, and continuity. More details on that project can be found in [12].

The vision of Space-Data Routers (Figure 4) is to forward data from space missions upon reception, whenever possible, directly to the interested parties (e.g. scientists, research institutes, etc.), utilizing a DTN overlay and applying CGR for routing decisions. The implemented version of CGR applies policy-based forwarding as an alternative to minimizing delivery latency. The contact plan and the forwarding criteria include the level-of-trust of each network node (for confidential data), the storage availability at each node, and a cost rate, enabling data to be forwarded over the lowest-cost route. Evaluation of the Space-Data Routers showed, among other conclusions, that CGR contributes to more efficient space data dissemination and has the potential to administer data confidentiality as well.

Operations on the International Space Station.

Pilot operation of the ION implementation of DTN on the International Space Station (ISS) officially began in July 2009: ION was installed in two commercial generic bioprocessing apparatus (CGBA) computers on ISS, where it was used to transmit science experiment telemetry to an experiment center at the University of Colorado (Boulder) continuously, via the Huntsville Operations Support Center at NASA's Marshall Space Flight Center in Alabama, for the next four years. The CGBA deployment on ISS was not topologically complex enough to fully exercise ION's implementation of CGR. However, the success of that pilot deployment convinced the ISS operations team that DTN would be a valuable permanent addition to the networking software infrastructure of the space station. Accordingly, two institutional DTN gateway nodes will be installed on ISS in 2015, serving both ISS operations and payload communications. The topology of the ISS DTN backbone will be as shown in Figure 5, including multiple potential cross-links in end-to-end paths.

The implementation of the DTN stack on the space station provides a variety of benefits, listed below:

- Enables payload developers (PDs) to automate operations and ensure science delivery with little regard for link or facility outages.
- Reduces the need for PD real-time support to access and downlink science data:
  - –DTN stores data during loss of signal (LOS) and automatically initiates transfer upon acquisition of signal (AOS).
  - –A download transfer can span Ku-Band AOS periods without any special scheduling or scripting.
  - –Reduces need for duplicate storage and extra retrieval actions.
- Reliable data transfer for ISS during LOS/AOS cycles:
  - –Automatic verification of bundle receipts, retransmissions reduced.
  - –When transmission errors occur, only the bundles that have errors are retransmitted, reducing the overall amount of retransmitted data, thus maximizing use of bandwidth.
- Allows PDs to use DTN protocols for their own applications (streaming, telemetry, etc.).
- Efficient use of downlink stream through DTN Quality of Service (QoS)/prioritization.
- Tolerance for high network latency (600ms delay is typical on JSL links).

Remarks

While the benefits noted in these experiments were provided only to satellite and space networks, the potential advantage in other networks where the DTN paradigm can be applied should be clear. In particular, the reliability guaranteed by DTNs can be very useful in mobile ad hoc networks (MANET) and, similarly in wireless sensor networks (WSN).

Conclusions and Future Developments

Routing in Delay-Tolerant Networks is a challenging problem, but practical solutions that apply principles underlying the design of the Internet are emerging. These solutions seem to offer the potential to support end-to-end data exchanges spanning a very wide range of communication environments.

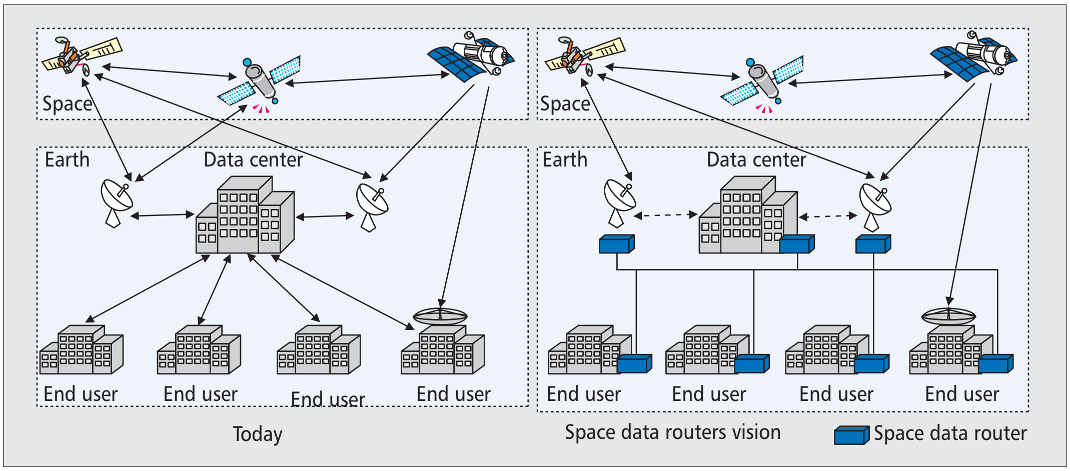

Figure 4. Space data dissemination today and as envisaged by the space data routers project [12].

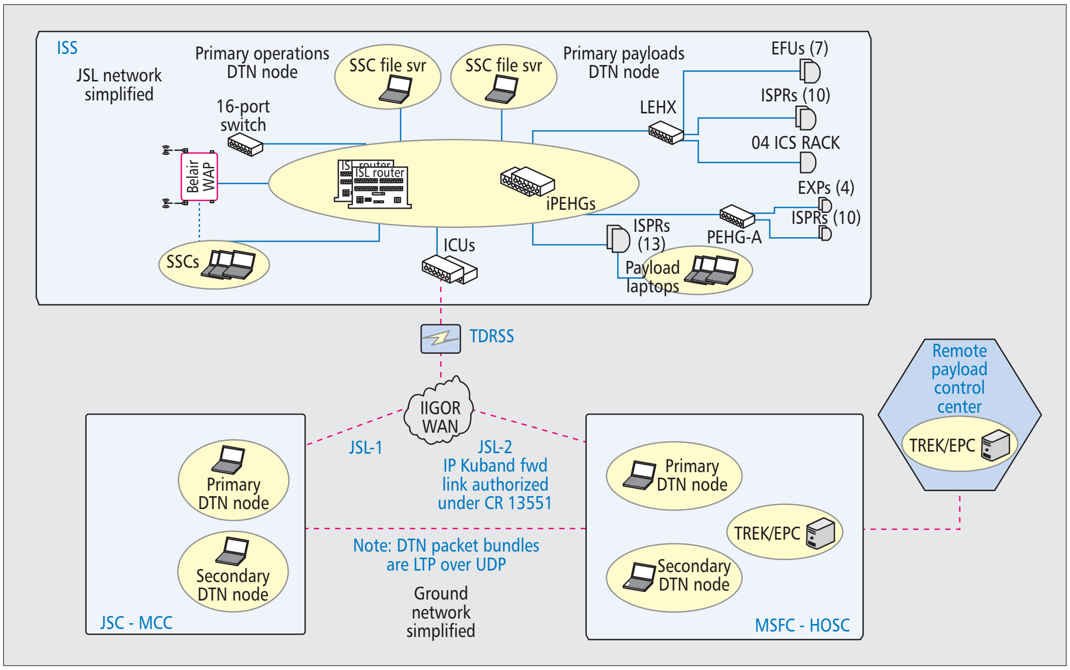

Figure 5. DTN on the international space station.

The applicability of CGR in the Ground segments of space missions was evaluated in "Space-Data Routers" (www.space-datarouters.eu), a European FP7 project that exploited the DTN architecture to improve the dissemination of Space mission data with respect to volume, timeliness, and continuity.

Experience gained in current and future deployments of Contact Graph Routing is vital to the future development of this technology. As noted earlier, future developments of CGR include the "Path Encoding CGR" and "Opportunistic CGR" approaches. In the former case the selected path within the space networks is memorized for validation and re-use by resource-constrained downstream nodes, while in the latter case a decomposition of the network is performed so as to apply a "divide and conquer" strategy.

Another point to be carefully addressed in the future development of CGR as well as other routing/forwarding approaches is processing efficiency. CGR was built for space exploration networks with scheduled communication opportunities, represented as a contact graph. Since CGR uses knowledge of future connectivity, the contact graph can grow rather large. Efficient processing approaches will be required to enable CGR to scale to anticipated future space network complexities.

Acknowledgments

This research was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Government sponsorship is acknowledged.

References

[1] Jet Propulsion Laboratory, California Institute of Technology, Exploration & Observational Systems, Deep Space Communications, Available: http://scienceandtechnology.jpl.nasa.gov/research/ResearchTopics/topicdetails/?ID=67

[2] M. Marchese, "Interplanetary and Pervasive Communications," *IEEE Aerospace and Electronic Systems*, vol. 26, Issue 2, Feb. 2011, pp. 12–18.

Routing in delay-tolerant networks is a challenging problem, but practical solutions that apply principles underlying the design of the Internet are emerging. These solutions seem to offer the potential to support end-to-end data exchanges spanning a very wide range of communication environments.

[3] V. Cerf *et al.*, Delay-Tolerant Networking Architecture, Internet RFC 4838, Apr. 2007, available: http://www.Rfc editor.org/rfc/rfc4838.txt.

[4] C. Caini *et al.*, "Delay- and Disruption-Tolerant Networking (DTN): An Alternative Solution for Future Satellite Networking Applications," Proceedings of the IEEE, Invited Paper, vol. 99, issue 11, July 2011, pp. 1980–97.

[5] K. Scott and S. Burleigh, Bundle Protocol Specification, Internet RFC 5050, Nov. 2007, available: http://www.rfc-editor. org/rfc/rfc5050.txt.

[6] S. Jain, K. Fall, and R. Patra, "Routing in a Delay Tolerant Network," *Proc. ACM SIGCOMM Portland*, Aug./Sep. 2004, pp. 145–57.

[7] S. Burleigh, Contact Graph Routing, Internet-Draft, July 2010, available: http://tools.ietf.org/html/draft-burleigh-dtnrg-cgr.

[8] N. Bezirgiannidis *et al.*, "Towards Flexibility and Accuracy in Space DTN Communications," *Proc.ACM CHANTS 2013*, Miami, FL, USA, Sept. 2013, pp. 43-48.

[9] N. Bezirgiannidis *et al.*,"Contact Graph Routing Enhancements for Delay Tolerant Space Communications," *Proc. ASMS 2014*, Livorno, Italy, Sept. 2014, pp. 17-23.

[10] E. Birrane, S. Burleigh, and N. Kasch, "Analysis of the Contact Graph Routing Algorithm: Bounding Interplanetary Paths," *Acta Astronautica*, vol. 75, June–July 2012, pp. 108–19.

[11] K. Suzuki *et al.*, "JAXA-NASA Interoperability Demonstration for Application of DTN under Simulated Rain Attenuation," *AIAA 2014, Proc. SpaceOps 2014 Conf*.

[12] M. Goetzelmann *et al.*, "Space Data Routers for the Exploitation of Space Data," *Proc. SpaceOps 2012 Conf*.

Biographies

GIUSEPPE ARANITI (araniti@unirc.it) is an assistant professor of telecommunications at the University Mediterranea of Reggio Calabria, Italy. From the same University he received the Laurea (2000) and the Ph.D. degree (2004) in electronic engineering. His major area of research includes enhanced wireless and satellite systems, traffic and radio resource management, multicast and broadcast services, device-to-device and machine type communications over 4G/5G cellular networks.

NIKOLAOS BEZIRGIANNIDIS (nbezirgi@ee.duth.gr) received a diploma in electrical and computer engineering from Aristotle University of Thessaloniki, Greece in 2006. He also received his M.Sc. degree in the scientific area of computer networks in the Department of Electrical and Computer Engineering, Democritus University of Thrace, Greece, in 2010. He is currently a Ph.D. candidate under advisor Prof. Vassilis Tsaoussidis in the same department.

EDWARD BIRRANE (Edward.Birrane@jhuapl.edu) is a senior embedded software engineer, principal investigator, and technical lead in a range of areas including embedded terrestrial telecommunications, space flight software, bioengineering, and national security space at the Johns Hopkins University Applied Physics Laboratory.

IGOR BISIO (igor.bisio@unige.it) received his "Laurea" degree in telecommunication engineering at the University of Genoa, Italy in 2002, and his Ph.D. degree in 2006. He is currently an assistant professor at the University of Genoa. He is the current Chair of the Satellite and Space Communications Technical Committe of the IEEE Communications Society. His main research activities concern resource allocation and management for satellite and space communication systems and signal processing over portable devices.

SCOTT BURLEIGH (Scott.Burleigh@jpl.nasa.gov) has been developing software at JPL since 1986. As a member of the Delay-Tolerant Networking (DTN) Research Group of the Internet Research Task Force, he co-authored the specifications for the DTN Bundle Protocol (RFC 5050) and the DTN Licklider Transmission Protocol for delay-tolerant ARQ (RFC 5326). He has developed implementations of these protocols that are designed to enable deployment of a delay-tolerant Solar System Internet.

CARLO CAINI (carlo.caini@unibo.it) received the master degree in electrical engineering "summa cum laude" from the University of Bologna in 1986. At present he is with the Department of Electrical, Electronic and Information Engineering (DEI) of the same University, as associate professor of telecommunications. His main research interest concerns satellite networking, transport protocols and delay tolerant networks.

MARIUS FELDMANN (marius.feldmann@tu-dresden.de) received his Ph.D. degree from the Technische Universität Dresden in 2011. He has worked for several national and international research projects, including a FP7 project funded by the European Union. His current professional focus lies upon entrepreneurship and he works as a lecturer in the domain of computer networks. His main research interests cover delay-tolerant networking as well as the history of Unix-like operating systems.

MARIO MARCHESE (mario.marchese@unige.it) received the "Laurea" degree (cum laude) in electronic engineering and the Ph.D. degree in telecommunications from the University of Genoa, Genoa, Italy, in 1992 and 1996, respectively. He is currently an associate professor with the DITEN Department, University of Genoa. His main research activity concerns satellite and radio networks, transport layer over satellite and wireless networks, and quality of service.

JOHN SEGUI (John.Segui@jpl.nasa.gov) is a senior technologist at Speedy Packets, Inc., and an instructor at Pasadena City College, CA, USA. He had been with the Jet Propulsion Laboratory and Riverside Community College. He received his bachelor degree in information and computer sciences at the University of California Irvine, and his master degree in computer science at the University of Southern California. His main research interest concerns are satellite networking and delay tolerant networks.

KIYOHISA SUZUKI (suzuki.kiyohisa@jaxa.jp) is currently an engineer in the Department of Consolidated Space Tracking and Data Acquisition of the Japan Aerospace Exploration Agency (JAXA), Japan. He is an active member of the Consultative Committee for Space Data Systems (CCSDS). His main research interest is in satellite networking, delay tolerant networks, and interoperability issues.