

Toward a Communications Satellite Network for Humanitarian Relief

Scott C. Burleigh

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California, USA
Scott.C.Burleigh@jpl.nasa.gov

Edward J. Birrane

The Johns Hopkins University Applied Physics
Laboratory
Laurel, Maryland, USA
Edward.Birrane@jhuapl.edu

Abstract—Since the introduction in 2008 of the “Ring Road” concept, proposing a communications satellite network designed to support disadvantaged populations, there have been a number of advances in the underlying technologies, CubeSat picosatellites and Delay-Tolerant Networking. We review the original Ring Road proposal, discuss relevant recent technological progress, and offer some tentative notes on projected cost and performance.

Keywords—communication satellite; CubeSat; picosatellite; Delay-Tolerant Networking

I. INTRODUCTION

The “Ring Road” proposal for a low-cost communications satellite network, based on the integration of CubeSat and Delay-Tolerant Networking (DTN) technologies, was introduced at the 59th International Astronautical Congress held in Glasgow, Scotland, in the fall of 2008. Since that time, advances in the technologies underlying Ring Road have helped to clarify the scope of this challenging concept. In this paper we review the original Ring Road proposal and then briefly discuss some relevant deployment and development progress achieved by the CubeSat and DTN communities over the past three years. We conclude with some very rough projections of cost and performance.

II. RING ROAD OVERVIEW

The proposed Ring Road system is a low-cost communications satellite network designed to provide high-latency but highly robust data interchange capability by using CubeSat¹ picosatellites as “data mules” in a network fabric established by Delay-Tolerant Networking.

A. Motivation

A discussion of the motivation for the Ring Road proposal and a survey of potential applications is presented in [1]. Briefly, the intent of the network design is to provide reliable electronic data transmission and reception service to most of Earth’s inhabited surface at such low cost that even the most economically disadvantaged can access the fundamental information resources of the Internet.

While the innately high latency of Ring Road communications rules out some kinds of network services

(Internet telephony, highly interactive Web browsing, and massive multi-player games, for example), supported applications would include:

- Warnings of disaster events
- Requests for relief services
- Relief worker consultation, reporting, and direction
- Search and rescue support in remote areas
- Disease control information
- Weather forecasts
- Fish and game migration data
- Commodity pricing
- Distance learning
- Acquisition of data from remote sensors
- Email
- Research queries

B. Key Features

1) Low Cost

Commercial communications satellite networks provide conversational telephone service and Internet access; as such, they must satisfy strict power and performance requirements supporting uninterrupted, low-latency, end-to-end connectivity to assets on Earth, as both telephony and the Internet protocols (IP) require continuous information exchange. The requirement for low-latency, continuous connectivity within the constellation is a significant cost driver, as it affects both the number of satellites in the constellation and the orbits of these satellites.

For example, positioning a satellite in a low-Earth orbit (LEO) is relatively inexpensive. However, LEO satellites are not always in view of ground communication stations on Earth. For this reason, LEO communication networks must comprise multiple satellites, raising complex space-to-space and space-to-ground connectivity issues. The establishment, management, and utilization of cross-links in a LEO constellation contributes substantially to the cost of network deployment and operations.

Conversely, geostationary (GEO) satellites always have connectivity to their ground stations and therefore require fewer satellites in the constellation, but launching satellites into GEO orbit is extremely expensive. Typically, GEO satellites defray their launch cost by incorporating multiple payloads, making them complex and time-consuming to build.

In contrast, the Ring Road network aims to provide reliable epistolary data transfer; as such, it need not ensure uninterrupted end-to-end connectivity to assets on Earth.

¹ Note that other types of small satellites could supplement or replace the CubeSats in this proposal. A CubeSat-based architecture is presented simply to make the argument as concrete as possible.

Removing the requirement for continuous connectivity relaxes the orbit/constellation size relationships that drive satellite system costs. By tolerating episodic contacts, Ring Road may comprise solely nadir-pointing CubeSats in low-Earth orbit: during the time that a satellite is out of the view of any ground station, it retains in-transit data in a queue in its own local storage medium, awaiting its next ground station overflight. The satellites are small and simple, built largely from inexpensive off-the-shelf components that are widely used by aerospace students, and launch cost per satellite is low.

2) *Incremental Deployment*

Because cross-links need not be maintained amongst Ring Road satellites, each satellite functions independently as a data transfer device. That is, there is no minimum topology required to enable network operation: Ring Road could provide data communication service – albeit with extremely high latency and at very low data rates – throughout its intended coverage area even if only a single satellite were in operation.

The incremental nature of the architecture offers two clear advantages:

a) *Scalability*

With increasing numbers of satellites, the total carrying capacity of the network increases. Because of the low cost of building and launching each satellite, deploying a very large network is relatively inexpensive. Perhaps more importantly, deployment of such a network need not be accomplished all at once, so no large initial investment is needed. Satellites can be deployed individually and opportunistically over a period of years as funding becomes

available.

b) *Reliability*

Because each satellite is an independent data transfer device, the failure of any single satellite has little impact on overall network performance. Network capacity and delivery timeliness would degrade gradually and gracefully with satellite outage, and failed satellites could be replaced quickly and at low cost.

C. *Architecture and Operation*

The central design principle of Ring Road is to use inexpensive DTN-enabled “courier” satellites, acting as “data mules”, to physically transport data between radio-equipped computers lacking Internet connectivity, termed “cold spots”, and radio-equipped computers that have Internet connections, termed “hot spots” – and thence to and from all the non-radio-equipped computers on the Internet. (The name “Ring Road” derives from this structure, “ring” connoting the orbital movement of the data while in storage aboard a courier and “road” connoting the “sneakernet” character of the DTN store-and-forward networking.) Fig. 1 illustrates this architecture.

Note that both the hot spots and (in most cases) the cold spots of Ring Road will be routers, functioning as gateways that forward traffic to and from other computers in IP-based networks – either local area networks in isolated locations or the Internet itself. That is, while the number of cold spots served directly by the Ring Road infrastructure will be limited by the distribution of courier overflight opportunities, the total number of users may be much larger.

Because the orbital movements of the couriers are

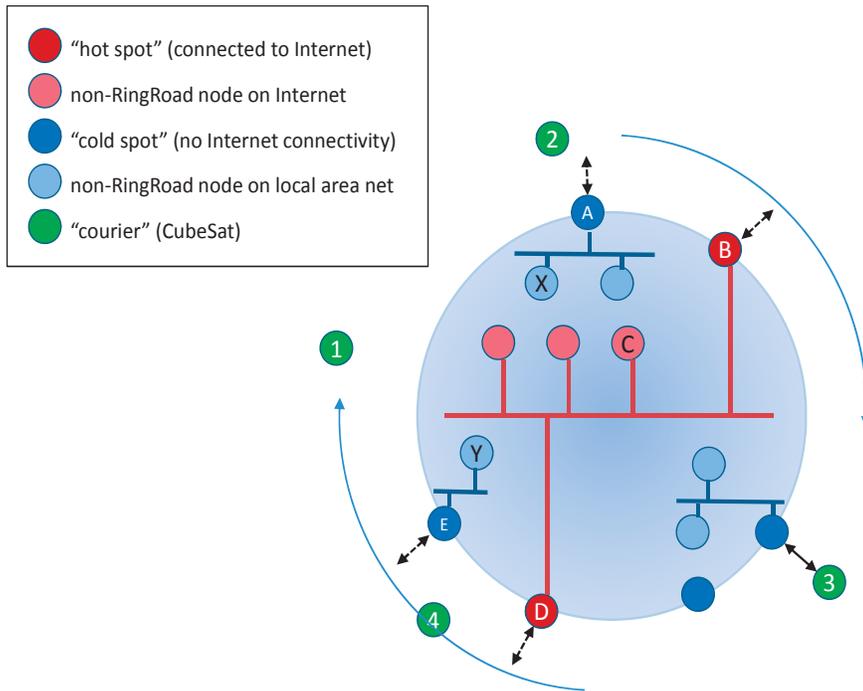


Figure 1 Ring Road Architecture

predictable, contact opportunities between couriers and hot and cold spots can be computed in advance. This enables optimal forwarding routes to be computed by means of the “contact graph routing” (CGR) procedures demonstrated during the four-week “Deep Impact Network” experiment conducted by NASA’s Jet Propulsion Laboratory in 2008 [2].

In Figure 1, for example, the isolated Ring Road user at computer “X” wishes to obtain data from a database server at “C”, which is accessible via the Internet. The user at “X” issues a query, destined for “C”, in a DTN bundle. The DTN protocol agent at “X” routes the bundle to the cold spot router at “A”. Router “A” briefly holds the query in local storage until the next Ring Road courier satellite – “2” – passes overhead, then transmits the bundle to “2”. Courier “2” retains the bundle in local storage while continuing on its orbit; when it passes over Internet-connected node “B”, it forwards the bundle to “B”, which in turn immediately relays the bundle via Internet to the server at “C”. The server responds with data packaged in a bundle and destined for node “X”. The DTN protocol agent at “C” determines that the fastest way to get the response back to “X” is to forward the bundle to courier “3”. (Although couriers “1” and “4” will be the next ones to be in view of router “A”, they won’t have an opportunity to receive any more data from the Internet prior to those overflights because there are no intervening hot spots.) So the bundle is immediately forwarded to “D”, the next Internet-connected DTN node that will be overflown by courier “3”. When “3” flies over “D” it receives the bundle destined for “X”, and when it subsequently flies over “A” it forwards the bundle to that router; router “A” then delivers to “X” the data from the database.

Alternatively, suppose the user at “X” wishes to send a message to another isolated Ring Road user at “Y”. The user at “X” issues the message to “Y” in a DTN bundle. The DTN protocol agent at “A” holds the message until courier satellite “2” passes overhead, then transmits the bundle to “2”. Courier “2” retains the bundle in local storage while continuing on its orbit; when it passes over Internet-connected node “B”, it forwards the bundle to “B”. The DTN protocol agent at “B” determines that the fastest way to convey the message to “Y” is to forward the bundle to courier “3”, the next courier that will visit “E”. (Courier “4” will have moved past “E” by that time.) So the bundle is immediately forwarded to node “D”, where it remains in storage pending the arrival of “3”. When “3” flies over “D” it receives the bundle destined for “Y”, and when it subsequently flies over “E” it forwards the message from “X”. The router at “E” then delivers the message to user “Y”.

D. Related Work

Students and faculty at Taylor University developed TU Sat 1, a communication CubeSat intended for launch in 2002 (<http://cse.taylor.edu/~physics/picosat>). TU Sat 1 included an email communications system capable of operating at a 115-kbps rate near 0.9 GHz. The goal was to demonstrate

low cost communication for remote villages in third world countries.

The satellite data services offered by ORBCOMM (<http://www.orbcomm.com>) are likewise much like email, storing data received while passing over one part of Earth and transmitting it later while passing over another part.

The US Army SMDC-ONE satellite launched on 8 December 2010 (<http://www.amsat.se/?p=7890>) was planned as the first satellite in a simulated tactical communications capability based on store-and-forward procedures, similar in concept to Ring Road.

The CASCADE payload of the Canadian CASSIOPE satellite (<http://mertensiana.phys.ucalgary.ca/cassiope.html>) will demonstrate a secure digital store-and-forward “courier” file delivery service.

Ring Road differs from this other work primarily in being based on standard DTN architecture rather than on an ad-hoc and/or proprietary store-and-forward mechanism. The features of this comprehensive architecture will enable high-volume transmission of a very wide range of data types (potentially even including streaming – though not real-time – video) with high reliability and security, at minimized end-to-end latency.

E. CubeSat Overview

The program to devise the inexpensive “picosatellites” called CubeSats began at California Polytechnic State University and Stanford University in 1999. CubeSats were originally conceived as a teaching tool, for educating aerospace engineers, but over the past decade dozens of CubeSat projects have been undertaken not only by universities but also by government agencies and some corporations [3].

A CubeSat is an assembly of from one to three standard bus structures, each one a 10-centimeter cube with mass of up to 1 kilogram (see Fig. 2).

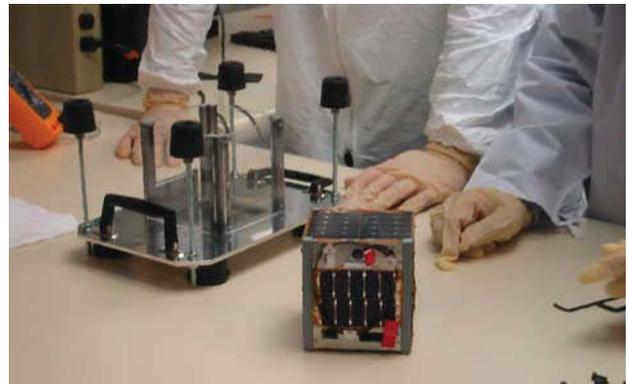


Figure 2 CubeSat at University of Tokyo (used by permission – <http://www.cubesat.org>)

Power is normally supplied by photovoltaic cells mounted on the exterior surfaces of the cube(s). Attitude control is typically passive, accomplished by bar magnets and hysteresis rods that simply align the satellite with the Earth’s own magnetic field. The flight computer is typically a capable x86-class single-board computer in the PC/104

form factor, and on-board storage may consist of as much as 2 MB of static RAM and up to 4 GB of flash memory.

Standardizing the satellite bus has enabled growth of a vigorous international satellite engineering community and the emergence of satellite “kit” vendors such as Pumpkin Space Systems. This in turn has lowered the cost of building a fully functional research picosatellite and launching it into Earth orbit to as little as \$52,000 [4].

F. DTN Overview

Delay-Tolerant Networking has been an active research field since the DTN Research Group of the Internet Research Task Force was established in 2002. The aim of the research was to develop a standard framework for automated network communication over “challenged” networks, i.e., networks built on communication links that are characterized by high signal propagation latency and/or lengthy lapses in connectivity. This work has culminated in the publication of several IRTF Requests for Comment including:

- RFC 5050, specifying the DTN Bundle Protocol (BP), a network protocol for data routing and forwarding over challenged networks;
- RFC 5326, specifying the Licklider Transmission Protocol (LTP), a protocol for ensuring data transmission reliability over the links of a BP-based network, or “dnnet”; and
- RFC 6257, specifying the Bundle Security Protocol (BSP), a set of extensions to BP that can ensure the authenticity, integrity, and confidentiality of the “bundles” of data carried by the network.

Note that security mechanisms were developed as integral elements of the DTN architecture from its inception. This capability is central to the viability of the Ring Road concept: the effective cost of the network service is sharply lowered by restricting traffic insertion to authorized users, so that bandwidth is not squandered in the conveyance of “spam”.

III. PROGRESS SINCE 2008

Since the presentation of the original Ring Road paper in 2008, the maturity of the CubeSat and DTN technologies on which it is based has increased in several encouraging ways.

A. CubeSat Developments

In February of 2010 the U.S. National Reconnaissance Office (NRO) contracted with Boeing Phantom Works for up to 50 triple-unit CubeSats. One of the CubeSats was intended for use in weather monitoring; the missions of the other satellites were classified.

On 11 July 2010 the Indian Polar Satellite Launch Vehicle (PSLV) deployed two student-built CubeSats: TIsat 1 and Studsat. TIsat 1 was a Swiss satellite designed and built by the students and staff of SUPSI-DTI while StudSat was designed and built by undergraduate students from across India with support from ISRO.

On 8 December 2010 the SpaceX launch of the Falcon 9 spacecraft included successful on-orbit deployment of two CubeSats, one of which was the US Army SMDC-ONE satellite discussed above.

On 15 February 2011 the National Aeronautics and Space Administration (NASA) announced the selection of 20 small satellite projects to fly as auxiliary cargo aboard rockets planned to launch in 2011 and 2012. The selected missions were initiated by a high school in Virginia, universities across the country, NASA field centers, and Department of Defense organizations.

On 5 August 2011 the NASA Space Operations Mission Directorate announced its third CubeSat Launch Initiative in anticipation of launch opportunities planned for 2012, 2013, and 2014.

B. DTN Developments

The Interplanetary Overlay Network (ION) implementation of the DTN protocols, in addition to operating successfully in deep space for the four weeks of the DINET experiment in 2008 [2], began continuous operation on-board the International Space Station in low-Earth orbit in July of 2009 [5].

ION’s implementation of bundle authentication was successfully tested in a network spanning multiple firewalled flight centers in March of 2010. Implementations of the BSP procedures for bundle confidentiality and integrity are currently being tested.

Multiple initiatives to enhance and optimize contact graph routing are in progress at NASA’s Jet Propulsion Laboratory (JPL), the Johns Hopkins Applied Physics Laboratory (APL), the University of Bologna, and the Japanese national space agency.

In addition, the first steps toward standardizing DTN network management procedures are now being taken by JPL, APL, Ohio University, and others.

IV. TENTATIVE PRICE AND PERFORMANCE PROJECTIONS

A detailed engineering design of Ring Road is beyond the scope of this paper and indeed beyond the resources currently available for development of the concept. To help clarify the proposal, some rough “back of the envelope” calculations are presented below.

A. Size of the constellation

Given that the goal of the design is to provide continuous coverage to most of the inhabited surface of Earth, we should be able to compute a rough estimate of target constellation size.

We propose that Ring Road satellites will operate in near-circular low-Earth orbit at an inclination of about 50 degrees, tentatively at an altitude of about 500 km.

At that altitude, the available radius of visibility is somewhat over 1000 km. (That is, each courier is in view of all points on the ground within a circle of at least 1000 km radius, centered directly below the satellite.) The satellites will be most distant from one another at the equator. Since the Earth’s circumference is about 40,000 km, we can ensure satellite visibility at every point on the equator by populating 10 orbit planes.

We propose a target deployment of 15 satellites per orbit plane, a total target constellation size of 150.

B. Network capacity

Satellites in orbit at the altitude proposed for Ring Road travel at about 7.8 km/sec, an orbital period of about 90 minutes, which equates to 16 orbits per day. At this speed, the length of time a Ring Road courier satellite will be in view of any single router during any single overflight is about 256 seconds.

S-band radio transceivers capable of transmission at 230 kbps are available for CubeSats at relatively low cost (as little as \$700 each, though more expensive equipment that is more radiation-tolerant may be preferred). The use of these transceivers would enable up to about 7.2 MB to be uploaded to the network during any single courier overflight. The maximum upload capacity of any single orbit by any single courier would be about 20 times this figure, or about 144 MB, so the total maximum possible rate of data insertion into the network would be 144 MB per day per satellite. For a 150-satellite network, this comes to 21 GB/day or about 6 MB/sec, 48 Mbps.

Certainly the actual net rate of transmission over the network will be far less, perhaps only half this figure, since the satellites will be over uninhabited areas of the Earth's surface much of the time. We suggest, though, that even a net 24 Mbps of continuous network service would prove valuable to disadvantaged populations.

C. Deployment cost

The detailed cost analysis given by Jos Heyman in [4], performed in October of 2009, concluded that a university could reasonably expect the cost of constructing a CubeSat to be about \$50,000. That analysis assumed a configuration that differs somewhat from what would be expected of a Ring Road courier: notably, a camera was included and the transceiver was assumed to be UHF/VHF rather than S-band. Moreover, we expect the costs of some components and services identified in that analysis have increased over the past two years, and we recognize that the durability of the network will depend heavily on the robustness of the satellites' electronic and mechanical systems.

At the same time, however, we expect that the costs of some miniaturized electronic components have dropped since 2009. In addition, the project initiation and definition cost identified in Mr. Heyman's analysis would be amortized across the entire Ring Road constellation, reducing the unit cost significantly. Finally, we would hope to achieve some economies of scale in both the construction and launch of Ring Road courier satellites, given the volume of the planned launch activity. Bearing all these factors in mind, we suggest that a unit assembly cost of about \$100,000 per courier is not too implausible. For a constellation of 150 satellites, total construction cost would then be \$15 million.

Launch cost per satellite is difficult to estimate at this time, but in view of the emergence of an active commercial space flight industry we guess that an average launch cost of \$200,000 for a group of three CubeSats might likewise be not too implausible. If so, the total launch cost for a 150-satellite constellation would be \$10 million.

This brings the cost of the entire constellation to about \$25 million, only about five times the cost of constructing (not launching) a single Iridium satellite in 1998.

The cost of deploying Ring Road hot spots and cold spots is relatively small compared to the cost of deploying the couriers, but it should not be overlooked. We would expect the processing capability of any modern laptop computer to be adequate for routing Ring Road traffic; including the cost of an attached S-band transceiver, we would anticipate that these routing devices could be deployed for \$2000 each, if not less. The number and distribution of hot spots and cold spots would undoubtedly change over time, and no sound basis for guessing at the total number springs to mind. Strictly for illustration purposes, let us assume the deployment of about six times as many ground nodes as satellite nodes: 1000 hot and cold spot devices, for a total of an additional \$2 million.

D. Service cost

Suppose we estimate a total network lifetime of 5 years. If we arbitrarily assume that operating the network costs \$1 million per year, then the total lifetime cost of Ring Road is \$32 million.

At 24 Mbps, the total lifetime volume of traffic carried by the network would be just over 225 TB. The average cost of transmission over that time would then be about \$.13 per MB.

By way of comparison, these are some sample commercial satellite network charges as of late 2009:

- ORBCOMM service is offered through resellers. Reseller SkyMate's "platinum plan" charges in 2009 were \$1.40 per 1000 characters or about \$1433.60 per MB. (ORBCOMM service is typically limited to very brief messages.)
- Iridium service likewise is offered through resellers. Reseller WCC offered a pre-paid Iridium service plan for 5000 minutes over two years for \$4994.99, about \$1.00 per minute. Since the data rate offered by Iridium equipment was about 2400 bps, which is about 1 MB per hour, this rate equated to a cost of \$60 per MB.
- INMARSAT reseller Tempest Telecom offered RBGAN service at rates ranging from \$5.25/MB to \$11.60/MB and BGAN service at rates ranging from \$3.59/MB.
- Tempest Telecom also offered prepaid Thuraya DSL service at rates as low as \$.35/MB, still nearly three times our projected cost of Ring Road service.

V. FUNDING

Probably the greatest obstacle to implementing the Ring Road proposal is not technical but practical: Ring Road is not a model for a profitable business, so financing its deployment will not be straightforward.

Because Ring Road is specifically intended to serve the world's disadvantaged populations, we anticipate that funding will have to be obtained principally from individual charitable donors and non-governmental relief organizations.

Here again, though, the incremental nature of Ring Road deployment may be advantageous.

With each added courier, hot spot, and cold spot, the capability of the network increases in measurable proportion, so deployment of the network as a whole lends itself easily to the contribution of a large number of relatively small donors. Organizations of moderate means and even affluent private donors could fund individual network nodes and have the satisfaction of knowing exactly which pieces of network equipment their generosity made possible.

VI. CONCLUSION

The Ring Road proposal for deployment of low-cost network communications infrastructure has the potential to make a significant contribution to humanitarian relief efforts worldwide. The technical challenges are substantial, but we believe they are well understood and manageable. The accelerating maturity of the underlying technologies continues to sustain our confidence in the viability of this concept.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the expert guidance of Aaron Q. Rogers, of the Johns Hopkins University Applied Physics Laboratory, in correcting the errors in the original text of section IV.

Some of the research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] C. Krupiarz, C. Belleme, D. Gherardi, and E. Birrane, "Using SmallSats and DTN for Communication in Developing Countries", Proc. International Astronautical Congress (IAC '08), International Astronautical Federation, 2008, IAC-08.B4.1.8.
- [2] J. Wyatt, S. Burleigh, R. Jones, L. Torgerson, and S. Wissler, "Disruption Tolerant Networking Flight Validation Experiment on NASA's EPOXI Mission", Proc. First International Conference on Advances in Satellite and Space Communications (SPACOMM '09), IEEE Computer Society, 2009, pp. 187-196, doi:10.1109/38.
- [3] A. Q. Rogers and R. A. Summers, "Creating Capable Nanosatellites for Critical Space Missions," Johns Hopkins APL Tech. Dig. 29(3), 2010, pp. 283-288.
- [4] J. Heyman, "FOCUS: CubeSats – a Costing + Pricing Challenge", SatMagazine, Satnews Publishers, Oct. 2009.
- [5] A. Jenkins, S. Kuzminsky, K. K. Gifford, R. L. Pitts and K. Nichols, "Delay/Disruption-Tolerant Networking: Flight Test Results from the International Space Station", IEEE Aerospace conference, March 2010.