A Communication Satellite Network for Infrastructure-Poor Countries: Delay-Tolerant Networking on Cubesats

Scott Burleigh/Greg Kazz

Jet Propulsion Laboratory, California Institute of Technology

14 January 2010

Copyright 2010 California Institute of Technology. Government sponsorship acknowledged.

Agenda

- Discussion of the “Ring Road” communication satellite concept, not to sell it to you, but:
  - As a vehicle for talking about the concepts of Delay-Tolerant Networking (DTN)
  - As an initial focus for potential further discussions on the use of DTN in CubeSats
  - Highlight the System Engineering Aspects Involved

Systems Engineering Functions*

**** Defined in JPL SE practices Document***

1. Systems Architecting (purpose determination, concept formulation, structuring, and certification of use)
2. Requirements Development and Allocation
3. Analyze and Characterize the Design (Examine the –ilities and Do Trade Studies)
4. Manage Technical Resources and Performance
5. Identify and Work the Interfaces
6. Verification and Validation (Have we built the system right vs have we built the right system ?)
7. Systems Engineering Process and Reviews
8. Risk Management (Identify, Assess and Mitigate Risk)
9. Manage and Control Requirements and the Design (ECRs, PFRs,CM)
10. Managing the SE Tasks (How to implement the SE functions)

As a System Engineer, you will be acting as a Project Leader "influencing" the system design by providing the "big picture view". You will be acting as a Task Manager "managing" the cost and schedule performance *of system elements or the system. You will be acting as the* Engineering Technical Authority ultimately responsible for certifying the flight readiness of the project.

Problem

- Data communication can greatly benefit people in rural areas of developing countries:
  - Weather forecasts, health emergency notices, commerce information, students' research, etc.
- But data communication in rural areas of developing countries is limited.
  - No wired infrastructure.
  - Commercial communication satellites services are expensive, on the order of $5-$10 per MB.

Proposed Solution

- A communication satellite constellation that fulfills these high-level Requirements:
  - Offers service to all points on the planet
  - Can be built incrementally over time, so that some benefit can be delivered early while financing and deployment are still in progress
  - Is so inexpensive that it can be built and operated with donated funds and managed as a non-profit
  - Will be operated for TBD years and have a planned decommissioning phase

Ring Road Operations Concept

- Use Cubesats as communication satellites, termed *couriers*.
- Couriers are nadir-pointed, communicating only with the ground stations they fly over.
  - Delay-tolerant networking (DTN) protocols at levels 3 and above. At level 2:
    - Prototype could be AX.25 over VHF.
    - For better performance, Proximity-1 over UHF.
  - No crosslinks. No station-keeping. Minimal power needed.
- Couriers function as “data mules”, conveying DTN *bundles* from Internet-disconnected *cold spots* to Internet-connected *hot spots* and vice versa.

Familiar ComSat Model

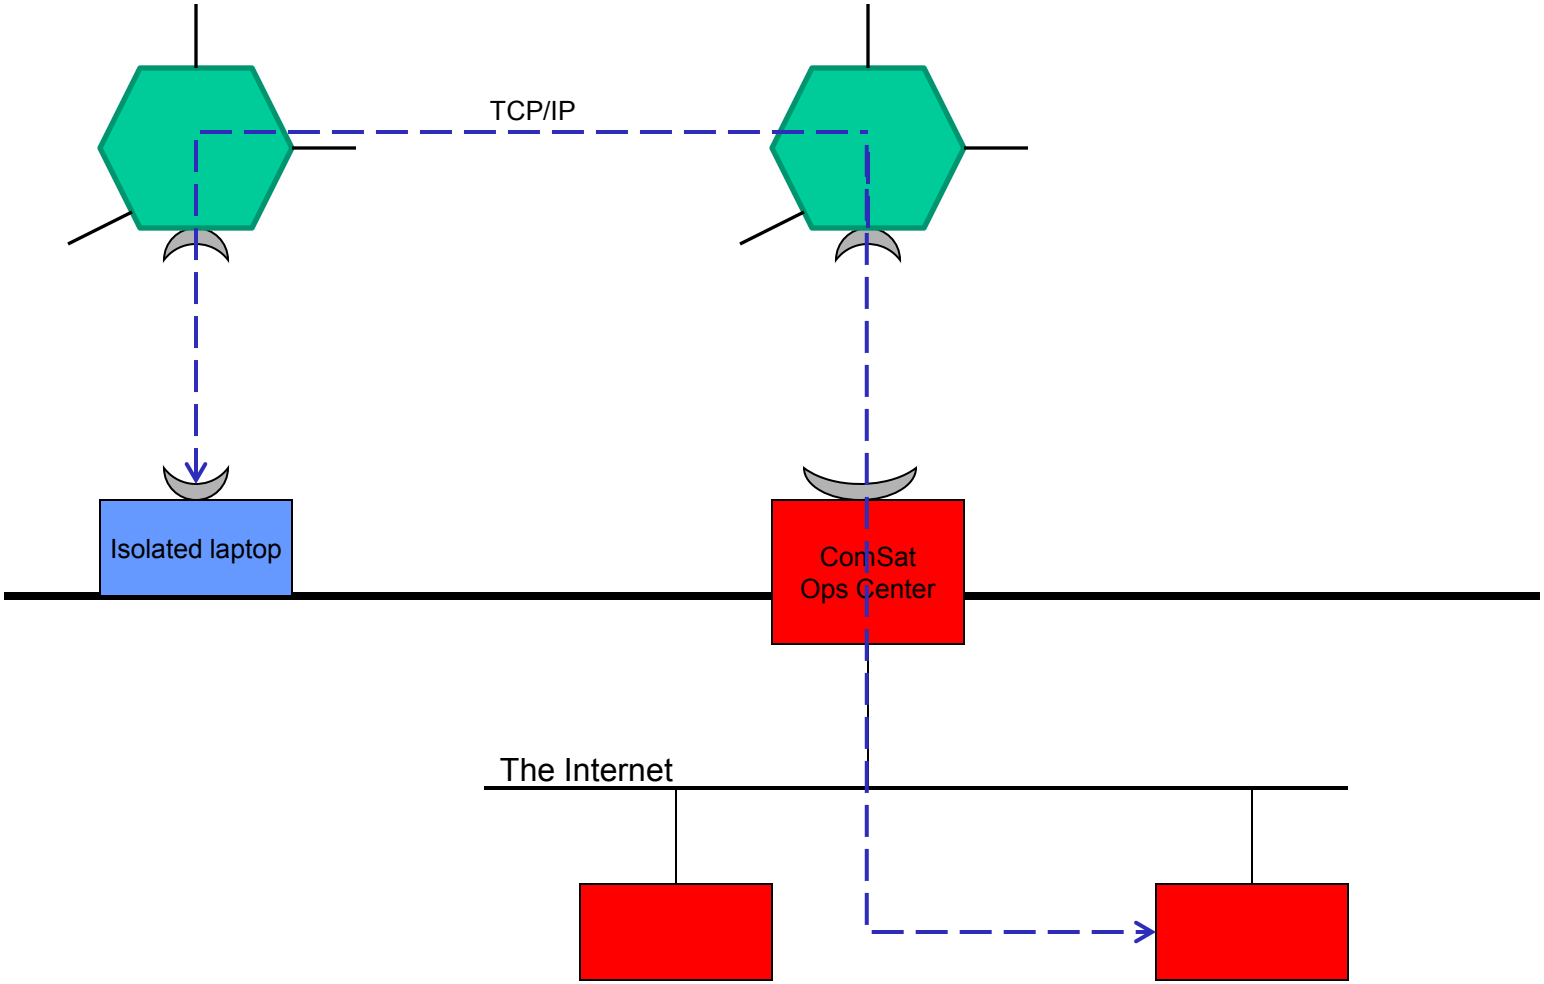

Ring Road Model

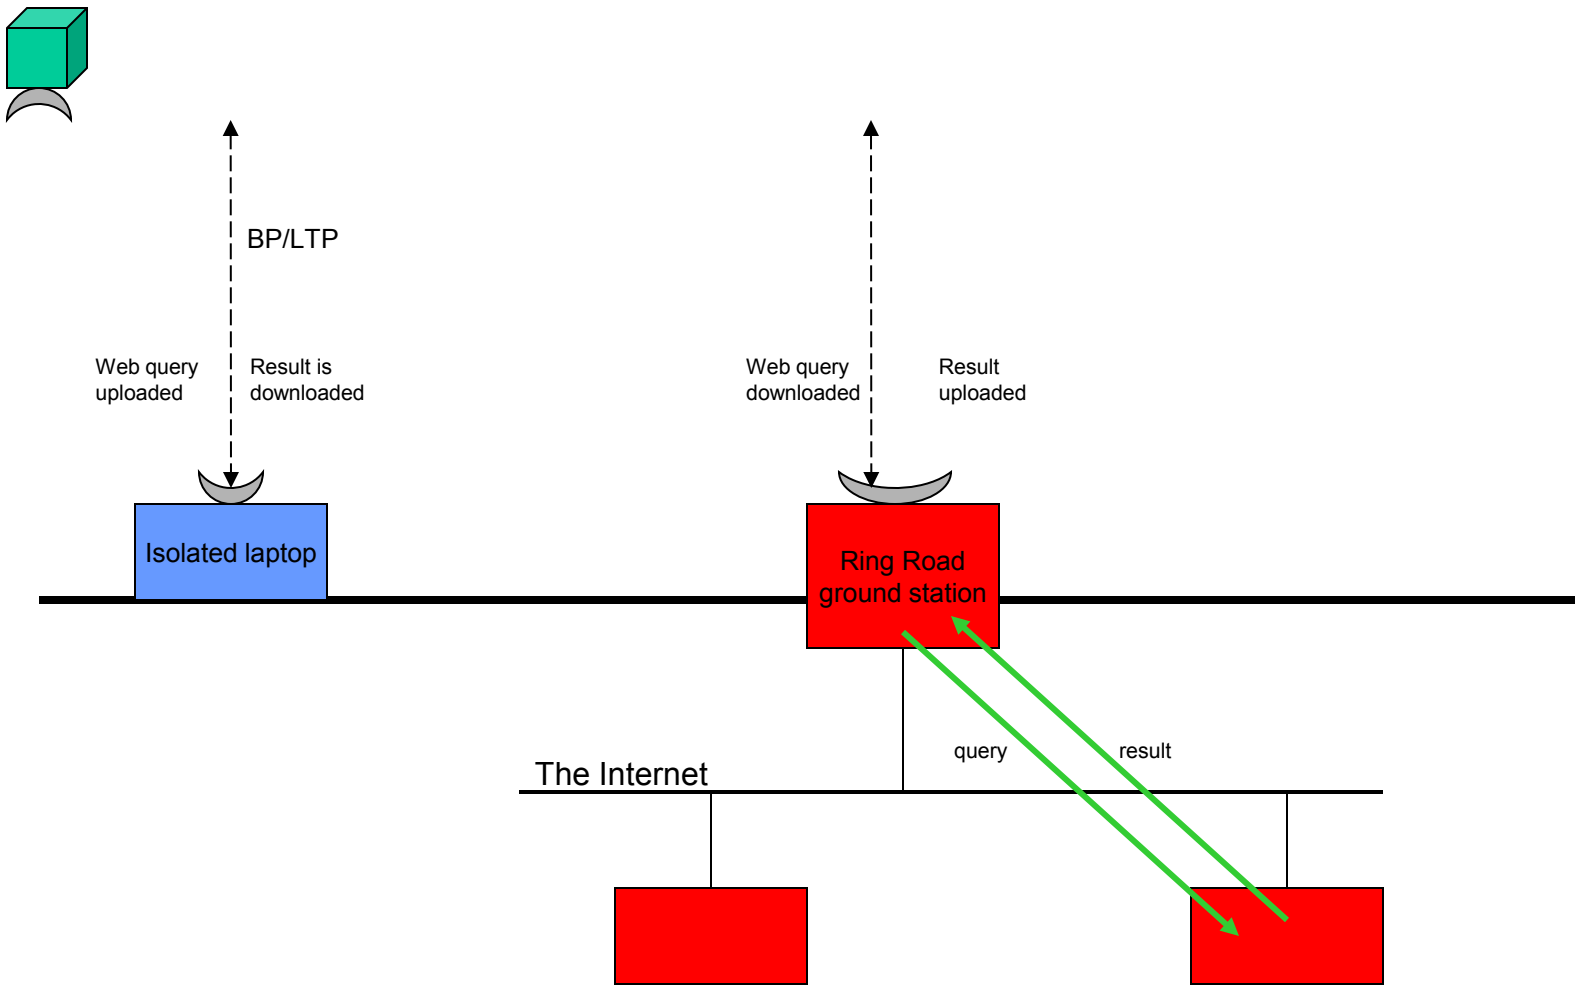

Scaling Up (1 of 2)
(Analyze & Characterize the Design)

- Couriers are in polar orbit, each one periodically flying over every point on the surface of Earth.
- As more “hot spots” are added, courier can offload data sooner, making room for more data. Network capacity increases.
- As more couriers are added, more data can be uploaded. Network capacity increases.
- As network capacity increases, more “cold spots” can be serviced.
- As more couriers and hot spots are added, round-trip latency is reduced.

Example 2

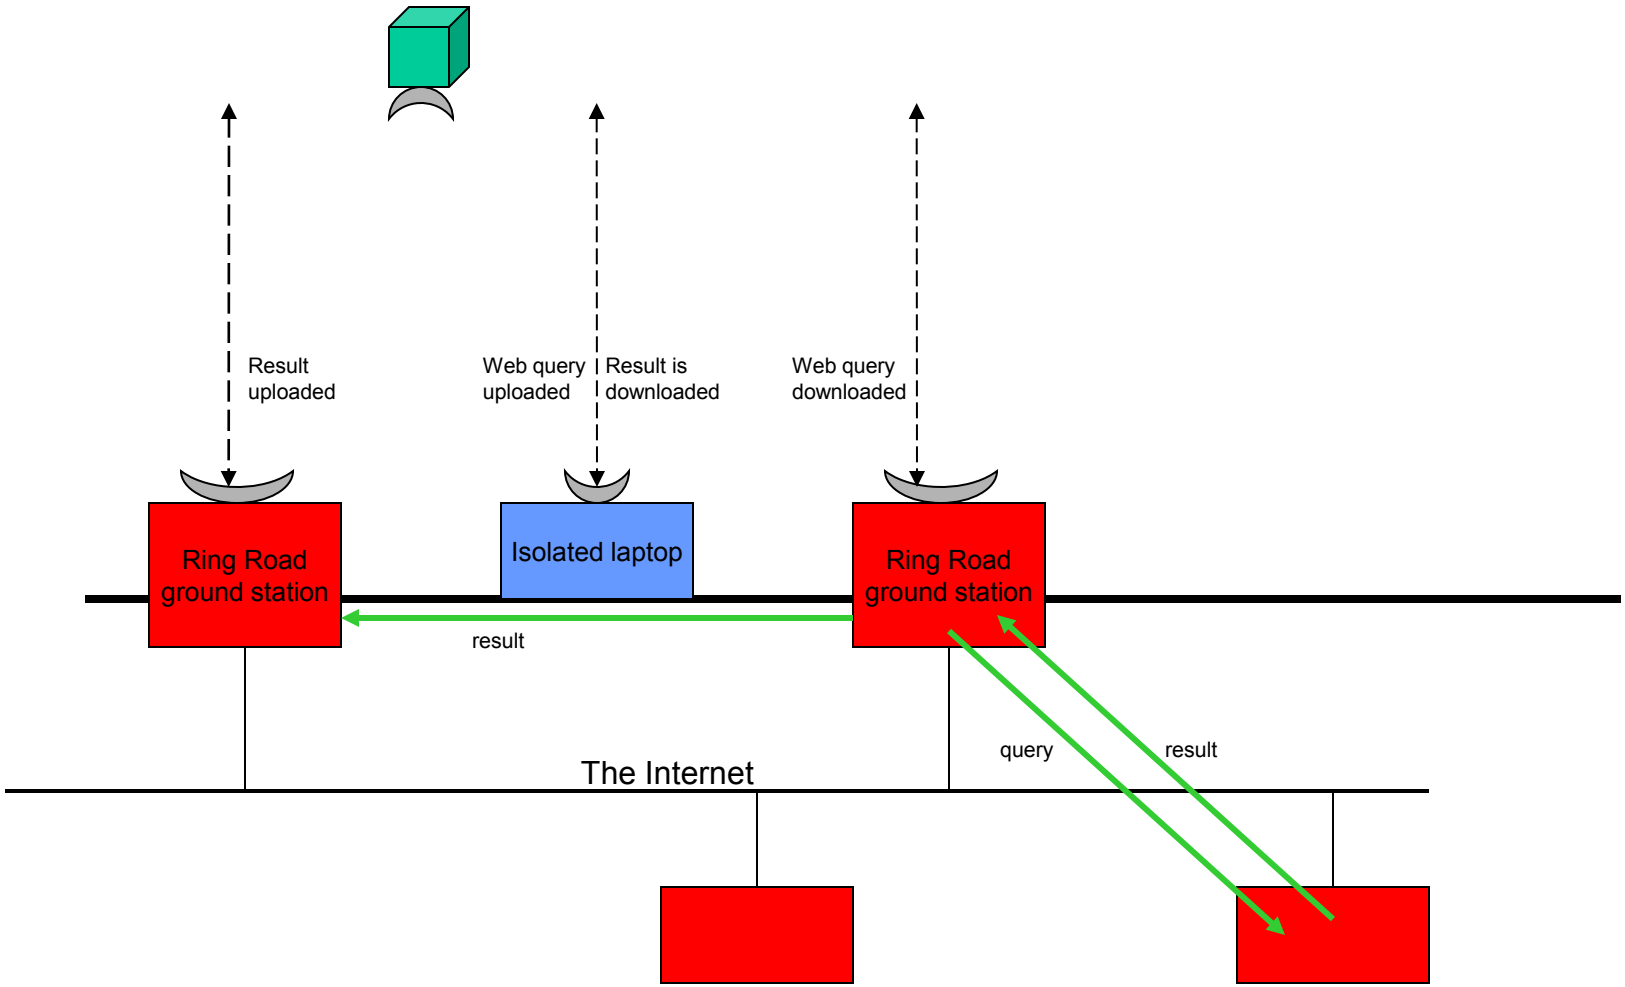
Result
uploaded
Web query
uploaded
Result is
downloaded
Web query
downloaded
Ring Road
ground station
Isolated laptop
Ring Road
ground station
result
query
result
The Internet

Scaling Up (2 of 2)

(Understand the system connections)

- The network can function – very slowly – with a single courier and a single hot spot. No need to launch the entire constellation before turning it on.
- Adding a hot spot requires just a laptop and a radio, a few $K.
- Adding a courier costs less than $100K.
- So NGOs, foundations, universities, corporations, and private individuals can incrementally fund the network over time.

The Technology

- Assuming triple-size Cubesats with Linux-capable processors and lots of memory.
  - Guessing no solid-state drives would be sufficiently cheap, small, and robust; true?
- DTN protocols on the couriers automate reliable communication with ground nodes and store-and-forward.
- DTN protocols on the hot spots automate routing for minimum round-trip latency.
  - Routing in the ground nodes minimizes processing load on the couriers.

Delay-Tolerant Networking: an Overview

Space Communications Today

- Communication opportunities are scheduled, based on orbit dynamics & operations plans.
- Transmission and reception events are individually scheduled, configured, and managed.
- Reliability over deep space links is by management: on loss of data, we command retransmission.
- More recently – MER – we have had managed forwarding through a relay point: we send commands to Odyssey or (until recently) MGS.

What's Wrong With That?

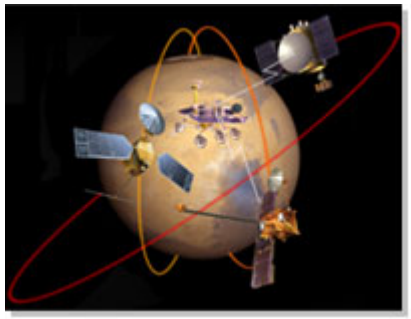

- This mission communications model has worked fine for over forty years; we've accomplished a lot of good science.
- But the status quo is:
  - Labor-intensive
    - Communication operations cost is a large fraction of the budget for each mission.
    - Risk of human error mandates mitigations that further increase cost.
  - Program-limiting
    - Cost and risk increase with the number of links between communicating entities.
    - As cross-links among spacecraft become common (Mars network, lunar exploration Constellation), cost and risk increases are non-linear with increase in the number of spacecraft.

2003/2004 Mars Relay Operations

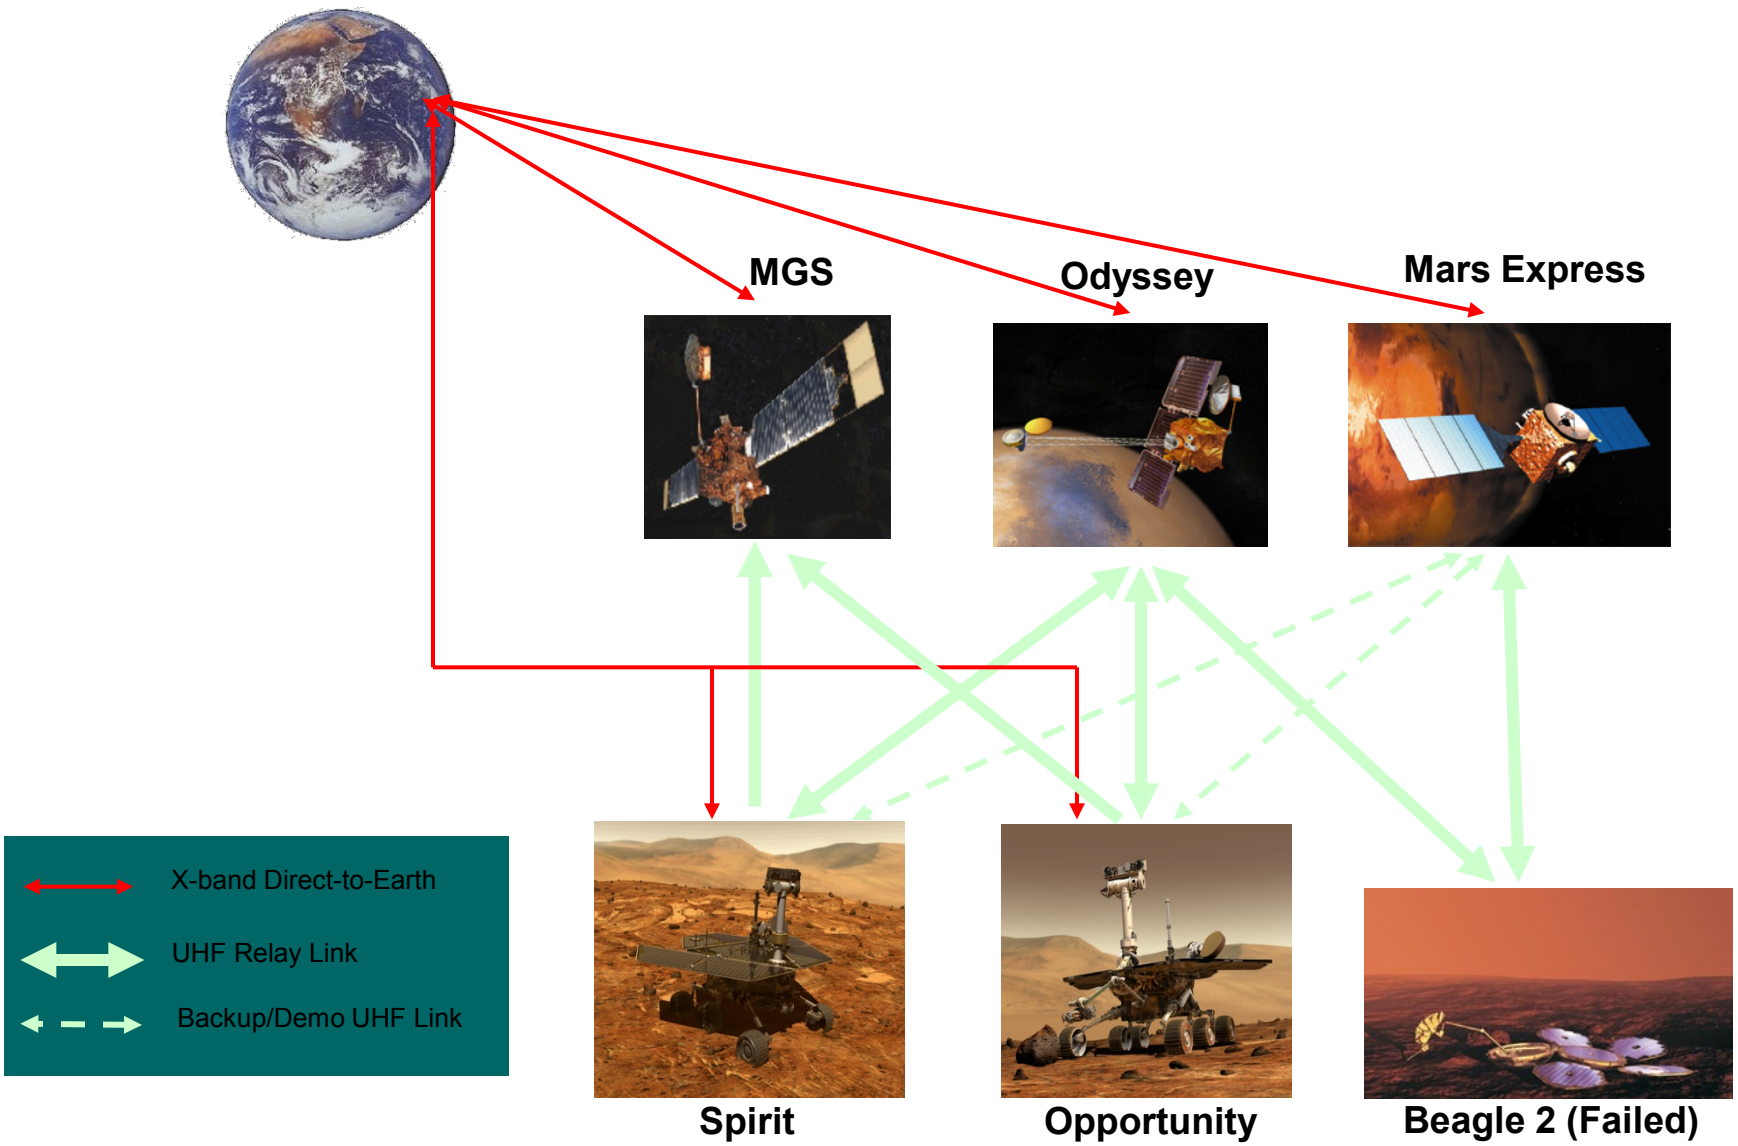

Proposed Future Mars Exploration Program

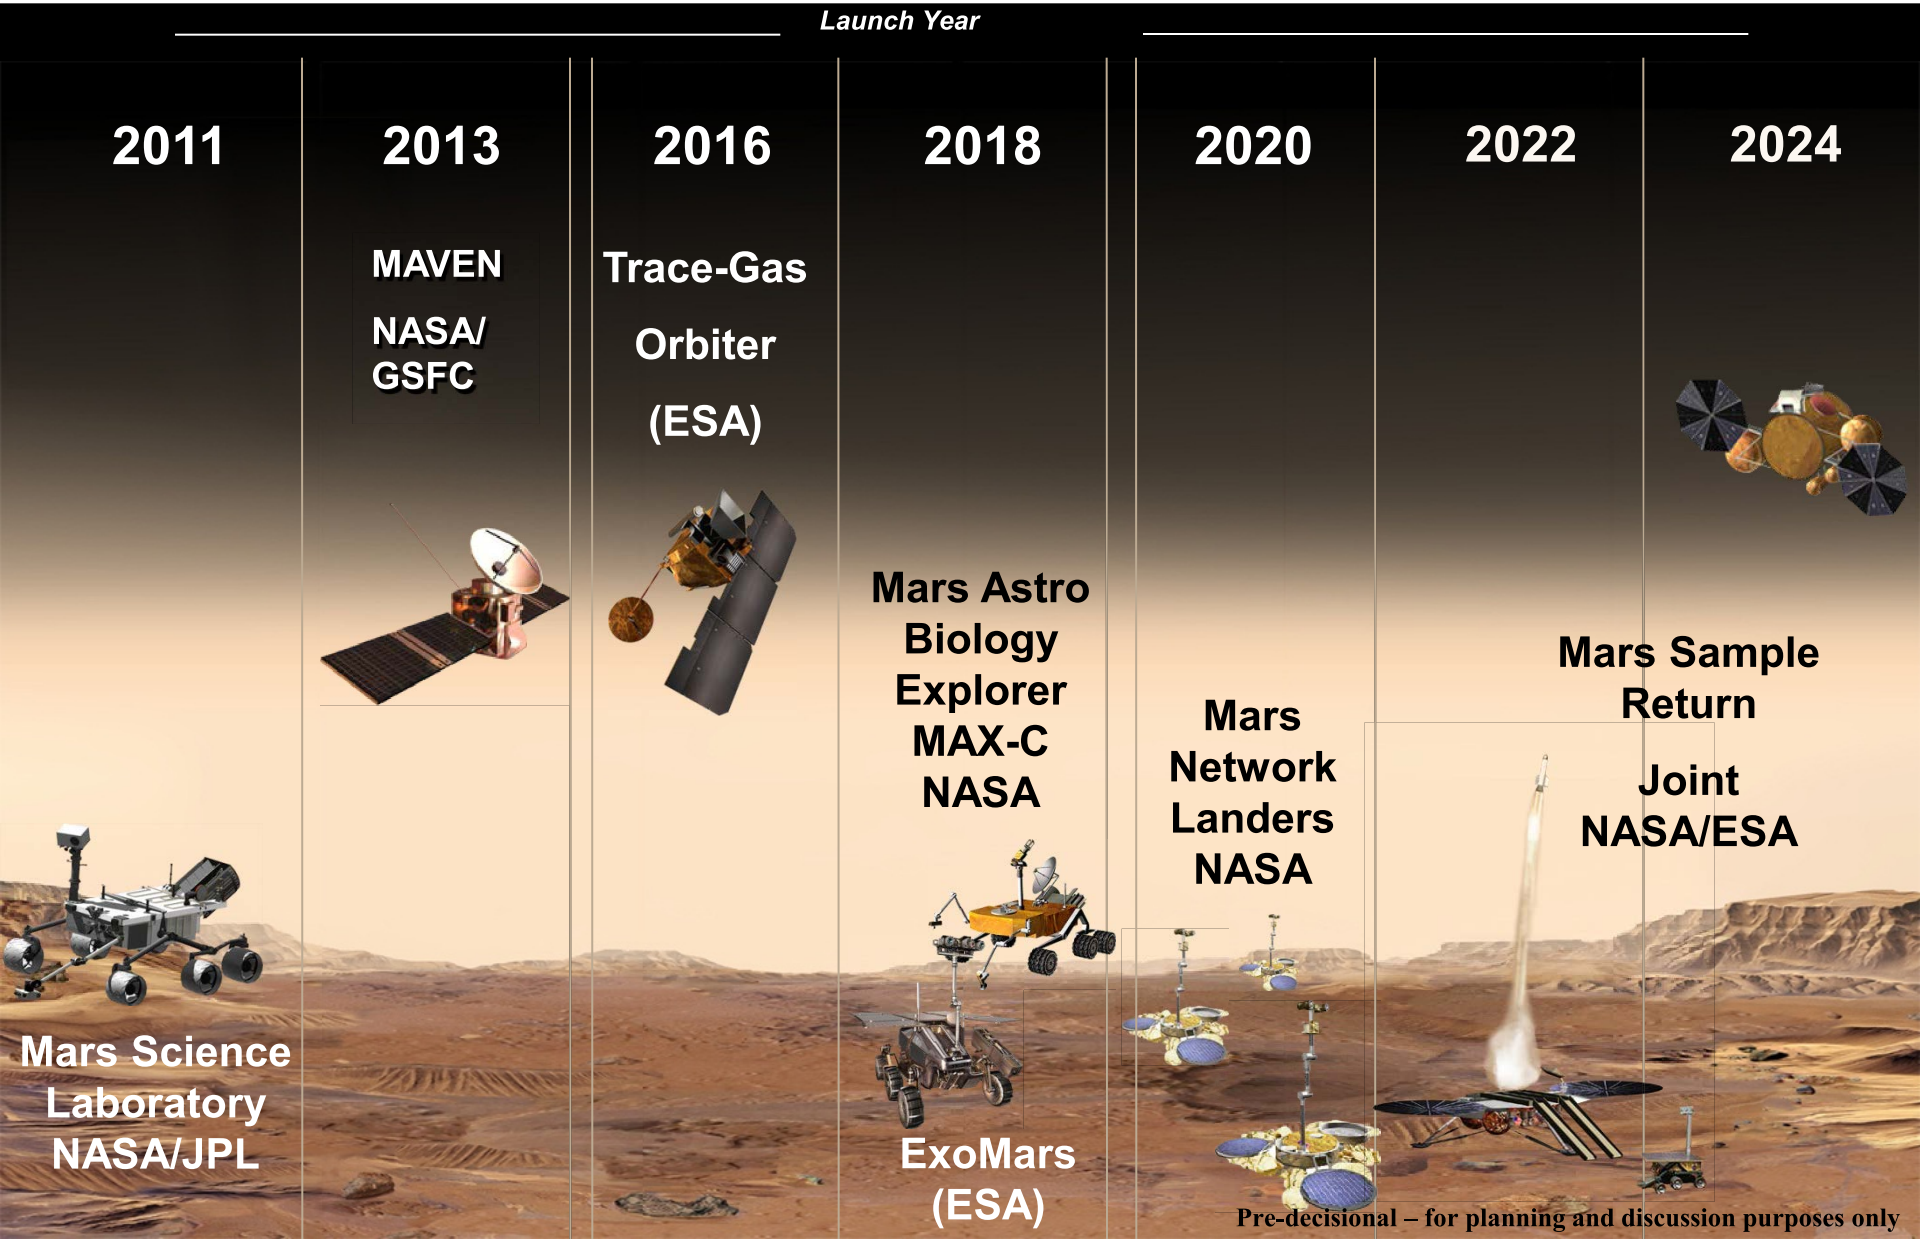

Pre-decisional – for planning and discussion purposes only

An Alternative

- The Internet has come to be widely used to conduct scientific investigations, for both science and engineering telemetry.
  - For example, the High-Performance Wireless Research and Education Network (HPWREN) in southern California.
    - Astronomy.
    - Ecology.
    - Geophysics.
- So why not use it for space science missions too?
  - Minimize cost (automation, COTS).
  - Minimize risk (huge installed base).

It Works Fine in Near-Earth Space

- Space Communication Protocol Standards (SCPS)
  - TCP options that improve performance on satellite links, where data loss is more often due to corruption than to congestion
  - international standard (CCSDS and ISO as well as DoD)
- Operating Missions as Nodes on the Internet (OMNI)
  - UoSAT-12, an HTTP server in orbit
  - CHIPSat, used Internet protocols on all communication links
  - CANDOS on STS-107, used mobile IP
- IP stack would likely also work well in <u>surface networks on other planets</u>.

So What's the Problem?

- Interplanetary space is a qualitatively different communication environment.
  - Internet, near-Earth, and planetary surface networks are all characterized by:
    - Very short distances between communicating nodes, therefore very **brief signal propagation delays** (up to about a second).
    - **Continuous end-to-end connectivity**. A network partition is treated as an anomaly and allowed to terminate communication.
  - Any network spanning interplanetary space would be characterized by:
    - Long distances between communicating nodes**, lengthy signal propagation delays** (e.g., 4-20 minutes from Earth to Mars).
    - **Routine network partitioning** due to lapses in connectivity on one or more links of the end-to-end path.

It's All About Delay

- Network disruption is, in essence, <u>highly variable</u> delay.
  - Case 1: continuous connectivity but client is, e.g., 56 million miles from server. *Response to query always arrives 10 min. after query is issued*.
  - Case 2: client and server are in adjacent offices but router is powered off for, e.g., 10 minutes. *Response to query, <u>on this occasion</u>, arrives 10 min. after query is issued*.
- Key effect of delay: reliable transmission of a given byte of data can take an arbitrarily long time.
  - Transmission can be lost due to corruption, N times.
  - NAK can be lost due to corruption, N times.
  - Disruption can delay transmission of NAK (or retransmission of data) by an arbitrarily long time.

Effects of Long and/or Variable Delay

- Connection establishment could take more time than entire communication opportunity.
  - So protocols must be connectionless.
- Transmission history can't be used to predict round-trip times.
  - So communication timeout interval computation must rely on link state information rather than timing statistics.
- End-to-end retransmission would reserve resources (retransmission buffer) at originator for entire duration of the transaction – possibly days or weeks.
  - So retransmission should be between relay points within the network rather than end-to-end: custody transfer.

Effects of Delay (cont'd)

- In-order stream delivery could be stuck for a long time, waiting for byte N to arrive before delivering byte N + 1.
  - So out-of-transmission-order delivery is needed – multiple concurrent transmissions.
  - So data must be structured in transmission blocks (e.g., files or messages) for concurrent retransmission – *not streams*.
- But reliable transmission of any single block can take an arbitrarily long time.
  - So any number of message transmissions might be in progress at the moment a computer is rebooted or power cycled.
  - So retransmission buffers should reside in non-volatile storage – not memory – to minimize risk of massive transmission failure.

Effects of Delay (cont'd)

- Propagation of topology changes in topology or DNS mapping could take an arbitrarily long time.
  - So route selection that is based on knowledge of the topology and/or DNS mapping of remote sectors of the network is more likely to be wrong than in a low-delay environment.
  - So either:
    - Changes in topology and DNS mapping must be limited, or
    - Routing must be based on some other information that is innately stable, such as endpoint <u>name</u>.
  - Limiting change in network configuration limits the usability of the network. It might work for a decade or so, but that would just be pushing the problem off until after we retire.
  - The long-term solution is **<u>late binding</u>** of endpoint names to network topology.

End-to-end IP in the Space Network

- For *unreliable* transmission, there are three problems:
  - Non-standard implementations would be needed to handle disconnection.
    - Network partition must not be treated as an enduring change in topology.
    - Network partition must cause outbound packets to be queued for future transmission rather than discarded.
  - The standard internal routing protocols use history-based timeouts to detect route failures. Routine transient partitioning would incorrectly cause route failure to be inferred and propagated, resulting in routing table errors.
  - Routing is based on IP address and known network topology, which might change too rapidly for routing protocols to track. Result would be routing errors.

End-to-end IP in the Space Network (2)

- For *reliable* transmission, there are additional problems:
  - Options are either (a) TCP or (b) UDP with reliability implemented in each application or middleware.
  - TCP isn't suitable.
    - It's based on connections, streaming, end-to-end retransmission, in-order delivery.
    - Retransmission buffers are in memory.
    - Timeout intervals are computed from transmission history.
    - Data loss is assumed to be due to congestion rather than corruption, so performance over noisy wireless links is degraded.
  - So the BGP external routing protocol standard, which uses TCP, is not suitable. So <u>none of the standard routing protocols work</u>.

End-to-end IP in the Space Network (3)

- More *reliable* transmission problems:
  - Without some single standard infrastructure for retransmission:
    - Each reliable application-layer protocol must reinvent retransmission – **additional cost and risk**.
    - **No standard for congestion control**.
  - Reliability from retransmission by the application is (like TCP) based on end-to-end retransmission. So:
    - No custody transfer. Network performance is degraded.
    - Timeout-triggered retransmission is excluded: end-to-end round trip time can't be predicted. Distributed Consensus problem.

Where Does That Leave Us?

- We could simply use IP anyway.
  - Omit routing protocols; just manage static routes.
  - Omit TCP, leave reliability to the applications and/or ops.
  - Omit congestion/flow control; just manage all data rates.
- But this would be functionally the same as status quo.
  - TCP-reliant Internet applications wouldn't work.
  - Would still be labor-intensive and program-limiting.
- Alternatively: develop **a new automated network architecture** that is tolerant of long and/or variable delay.
  - TCP-reliant Internet applications still won't work, but in some cases we can proxy them into the new infrastructure.
  - Reduce cost and risk: automate network functions, automate retransmission, integrate easily with Internet.

Delay-Tolerant Networking (DTN)

- An overlay network.
  - DTN “bundle protocol” (BP) is to IP as IP is to Ethernet.
  - A TCP connection within an IP-based network may be one “link” of a DTN end-to-end data path; a deep-space R/F transmission may be another.
- Reliability is achieved by retransmission between relay points within the network, not end-to-end retransmission.
- Route computation has temporal as well as topological elements, e.g., a schedule of planned contacts.
- Forwarding at router is automatic but not necessarily immediate: store-and-forward rather than “bent pipe”.
- Contain DOS attacks: reciprocal inter-node suspicion.

DTN Stack Elements

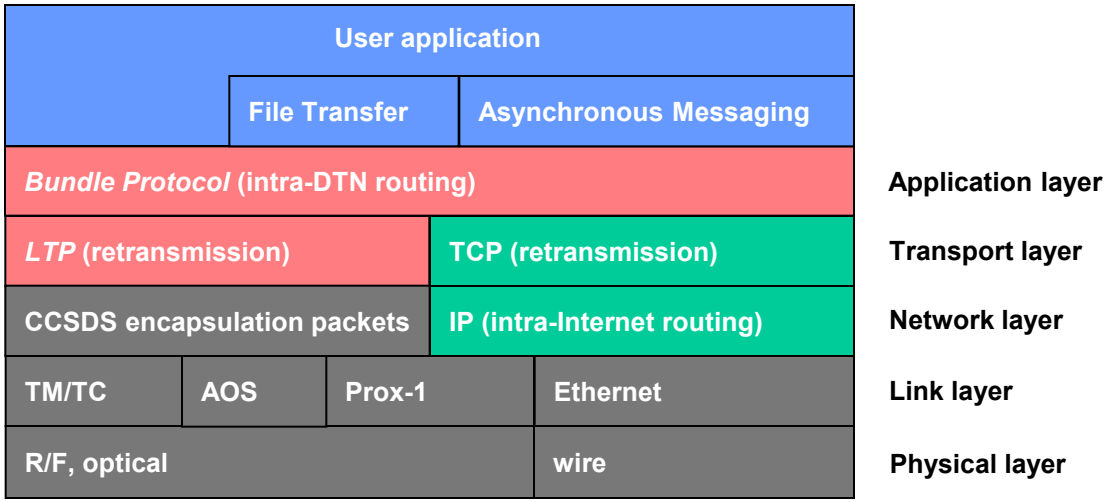

An End-to-End Path

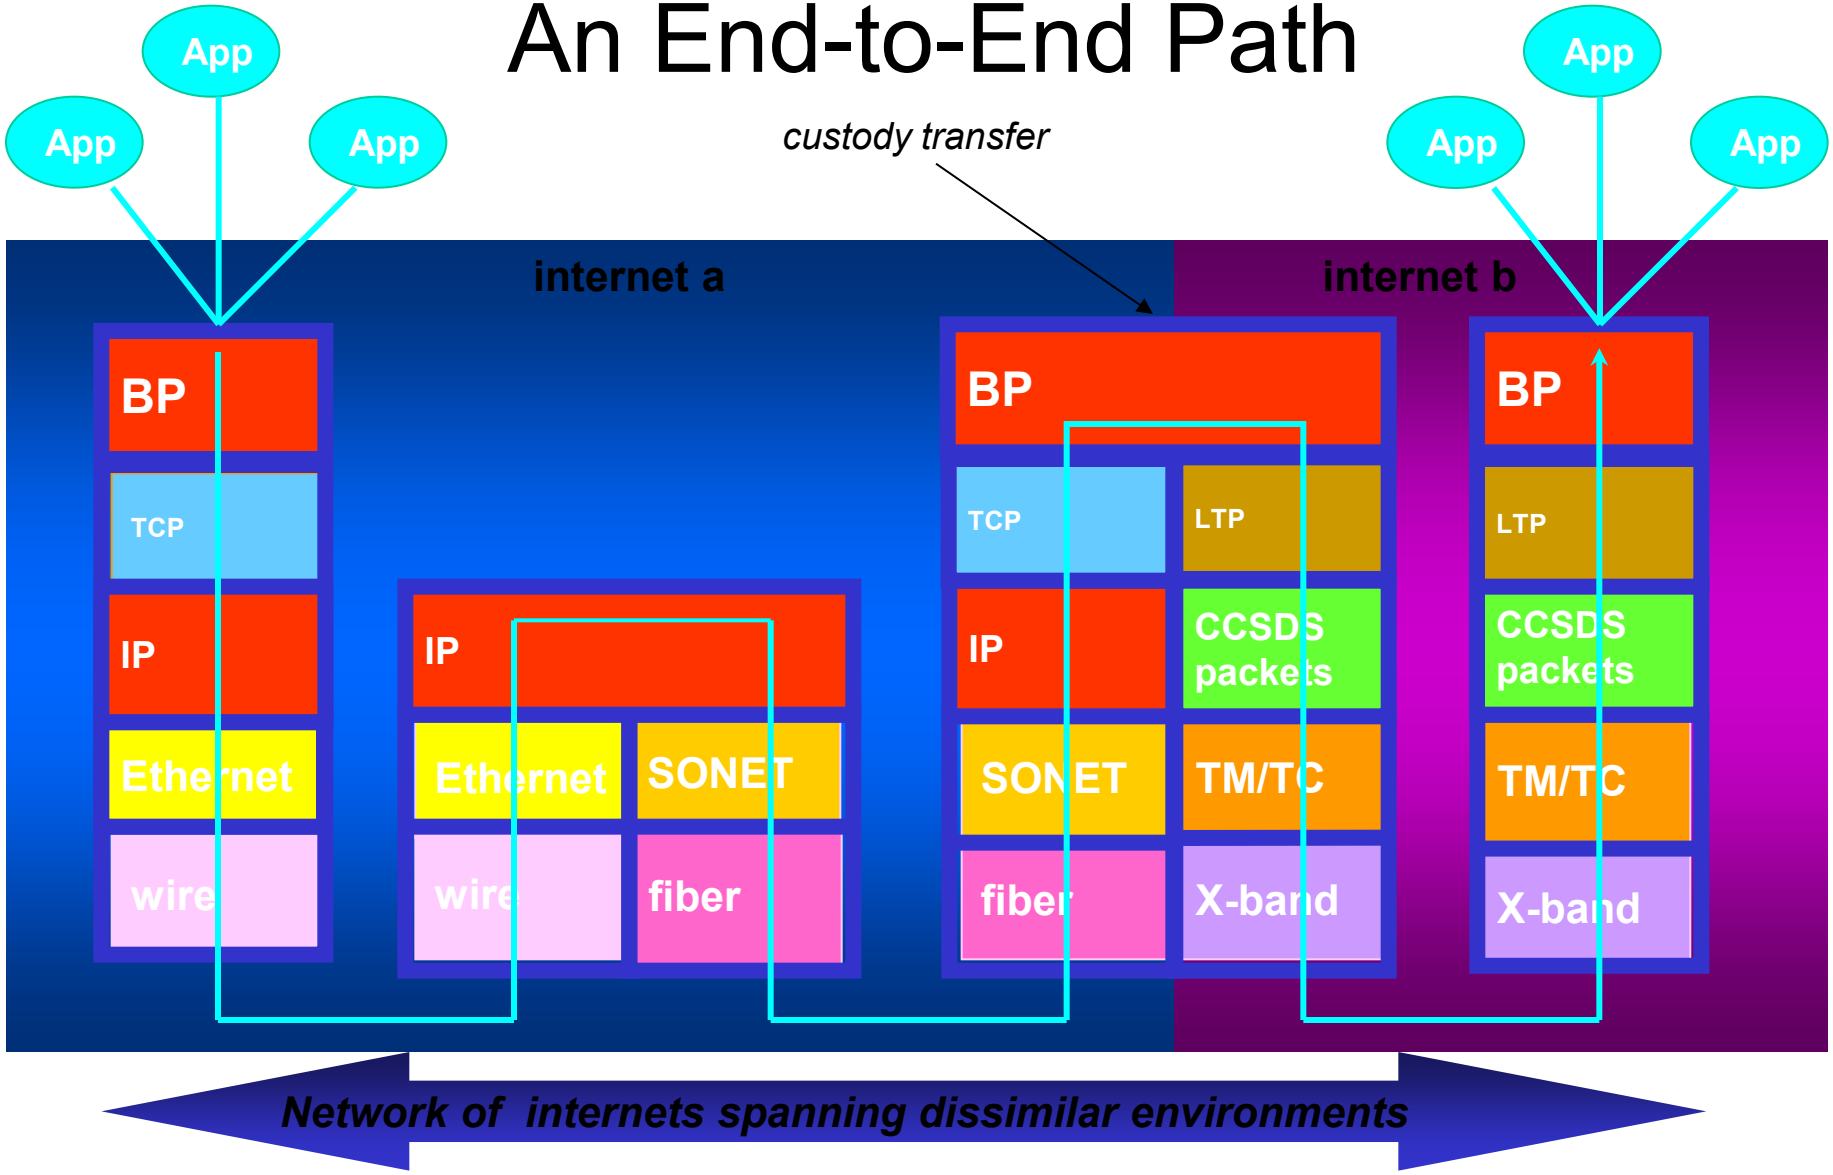

DTN Operations in space

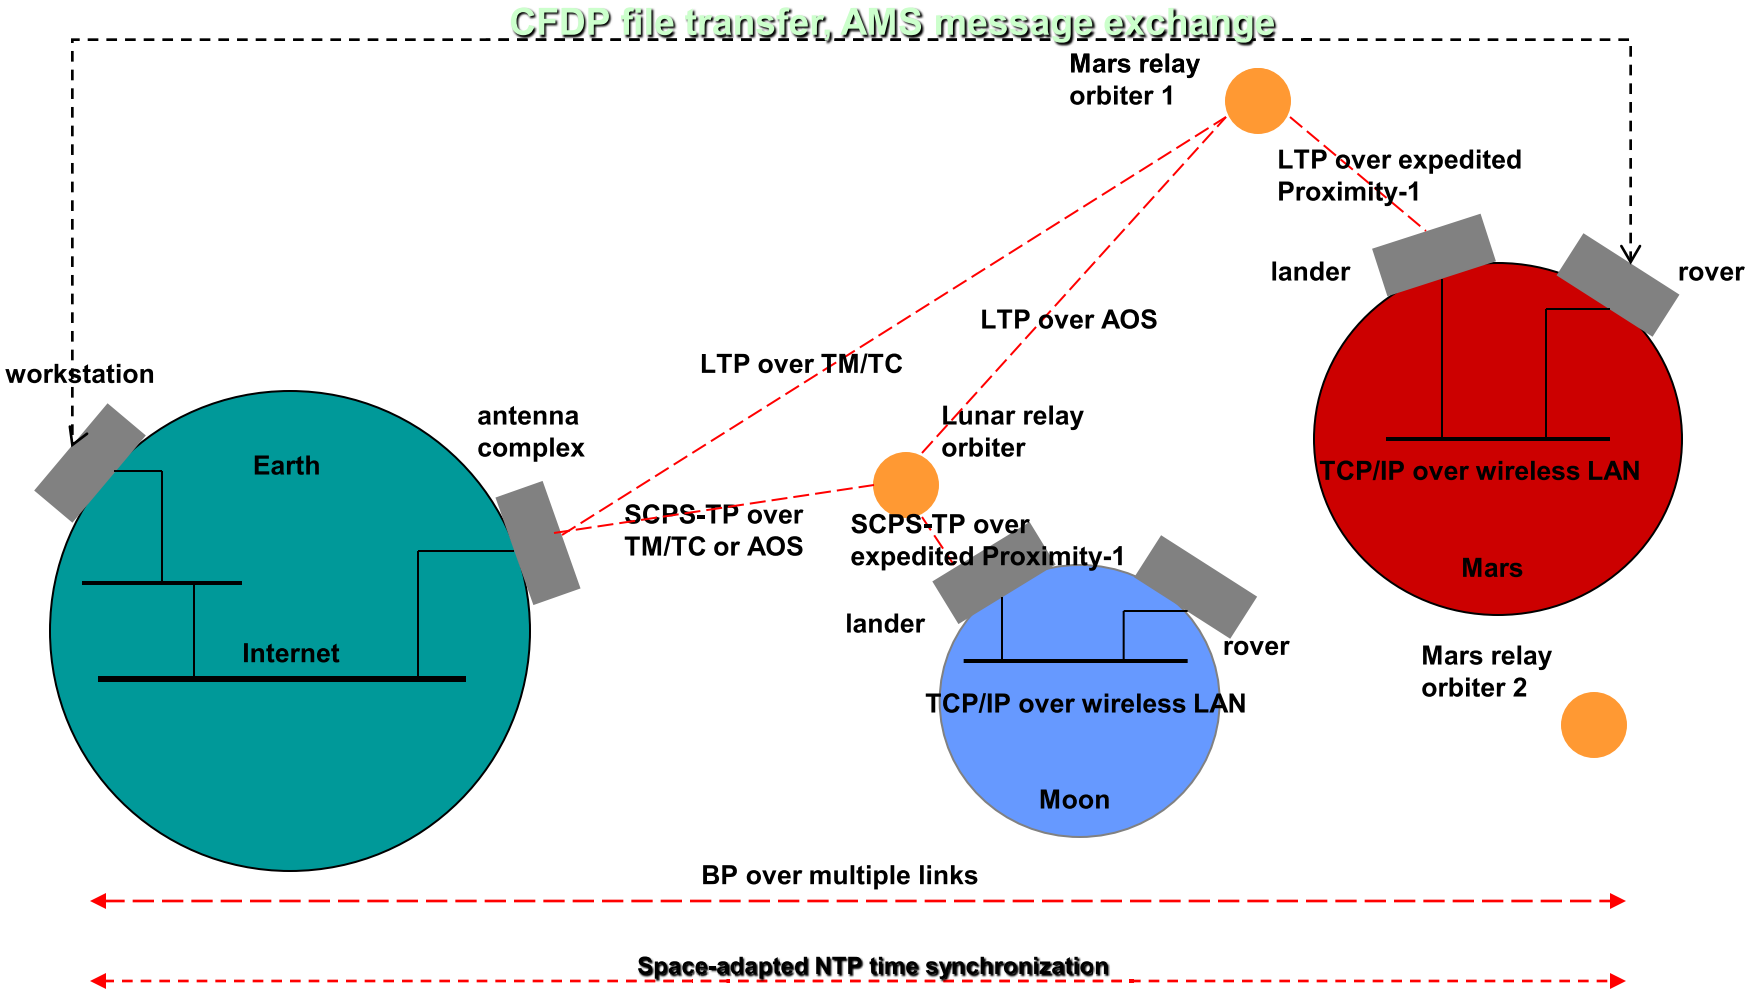

Other Applications of DTN

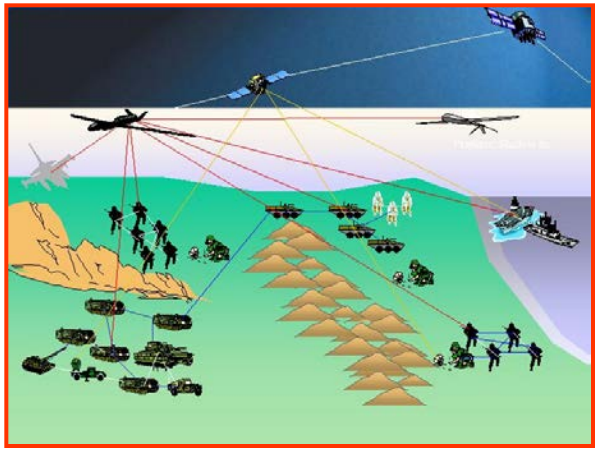

- Serving “challenged” regions of Internet topology
  - Mobile nodes: frequent disconnection, rapidly changing topology
  - Geographically remote nodes: relatively stable topology but infrequent connection
  - Undersea communication: long signal propagation delays, frequent disconnection
- Mobile tactical military communications
  - Frequent disconnection, rapidly changing topology

DTN Current Status

- Specifications and documentation
  - DTN architecture is RFC 4838
  - Bundle Protocol is RFC 5050
  - Licklider Transmission Protocol is RFC 5326
- Implementations
  - BP implementations
    - DTN2: open source reference implementation (Intel, UC Berkeley)
    - ION: designed for space flight (JPL)
    - At least three others are known
  - LTP implementations
    - Reference implementation in Java (Ohio University)
    - ION: designed for space flight (JPL)
    - Research implementation in C++ (Trinity College, Dublin)

Demonstrations

- Tactical military network demonstration at Ft. A. P. Hill in November of 2007.
- Demonstration on-board the EPOXI spacecraft, about 70 light seconds from Earth, in October-November of 2008.
- Currently operating in two instrument computers on-board the International Space Station.

Remaining Research Issues

- Route computation algorithms
  - Very different types of contacts
    - Scheduled
    - Opportunistic
    - Predicted
  - Traditional metrics (distance vector, link state) don't work.
    - They don't take timing into account: a two-hop path available in 10 minutes may be better than a one-hop path available tomorrow.
    - Topology may change too rapidly for protocols to track.
- Congestion control
  - TCP congestion window is end-to-end, may not reduce data injection rate at source until congestion collapse has already occurred.

Jet Propulsion Laboratory
California Institute of Technology

Career Opportunities

JPL is actively recruiting!!

We welcome new graduates and co-op students majoring in:

Computer Science/Computer Engineering
Electrical Engineering
Aerospace/Aeronautical Engineering
Mechanical Engineering
Physics/Mathematics

Opportunities available for all degree levels (B.S., M.S., Ph.D)

MINIMUM GPA OF 3.0 IS REQUIRED FOR ALL OPPORTUNITIES!!!

We also offer SUMMER INTERNSHIPS at JPL.
If you are interested in a summer internship or other types of opportunities, please visit our website at: http://careerlaunch.jpl.nasa.gov
JPL will be at Society of Women Engineers (SWE) Evening at CALPOLY on Jan 21, 2010