Deep Space Internet

BTEC Plenary Talk

Long Beach, CA
07 February 2006

Scott Burleigh
Systems Engineering Section
Jet Propulsion Laboratory, California Institute of Technology
818.393.3353
Scott.Burleigh@jpl.nasa.gov

Deep Space Communications Today

- Communication opportunities are scheduled, based on orbit dynamics & operations plans.
- Transmission initiation is manual, per schedule.
- Transmission direction is manual: point antenna, start transmitting when the right spacecraft is listening.
- Retransmission is manual: on loss of data, command repeat.
- More recently (MER), manual forwarding through relay point: command to Odyssey or MGS.

What's Wrong With That?

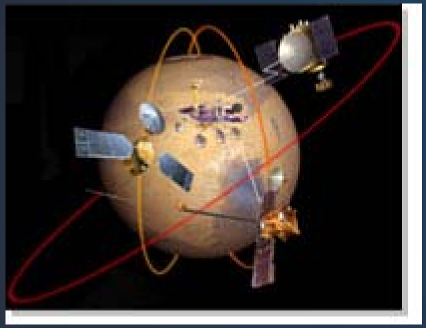

- This mission communications model has worked fine for over forty years; we've done a lot of good science.
- But the status quo is:
  - Labor-intensive
    - Communication operations cost is a large fraction of the budget for each mission.
    - Risk of human error mandates mitigations that further 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.

An Alternative

- The Internet has come to be widely used to conduct scientific investigations, for both science and engineering telemetry.
- So why not use it for deep space science missions too?
  - Minimize cost (automation, COTS).
  - Minimize risk (huge installed base).

High-performance Wireless Research and Education Network (HPWREN)

11 GHz link to Mt. Soledad

California Wolf Center

Optical Fiber Infrasound Sensor

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
- 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 cislunar space and in surface networks on other planets.

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 lapse in connectivity on any single link 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., 8-20 minutes from Earth to Mars).
    - Routine lapses in connectivity on all links of end-to-end path.

It’s All About Delay

- Network disruption is, essentially, unpredictable delay.
  - Case 1: continuous connectivity but client is 56 million miles from server. *Response to query arrives 10 min. after query is issued*.
  - Case 2: client and server are in adjacent offices but router is powered off for 10 minutes. *Response to query 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., 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.

Interplanetary IP – the Bottom Line

- None of these effects preclude the use of the IP network protocol (IP datagram transmission) itself.
- But:
  - TCP isn't suitable.
    - Connections, streaming, end-to-end retransmission, in-order delivery.
    - Retransmission buffers are in memory.
    - Timeout intervals are computed from transmission history.
  - The BGP external routing protocol uses TCP, so it's not suitable.
  - Internal routing protocols use history-based timeouts to detect route failures, so routine loss and re-establishment of connectivity would incorrectly cause route failure to be inferred and propagated to routing tables. Not suitable.
- The off-the-shelf IP stack doesn't work for deep space.

A Discursive Aside

- Tolerance of long and/or variable delay is what distinguishes postal, epistolary communications from telephonic, conversational communications.
- The postal model is more general: you can always pass notes to the person sitting next to you, but you can't always talk to your cousin in Burma.
- Nowadays we use cell phones to do both (text messaging).
- The Internet architecture is innately telephonic.

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.
- 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 arbitrary 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 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 for Deep Space

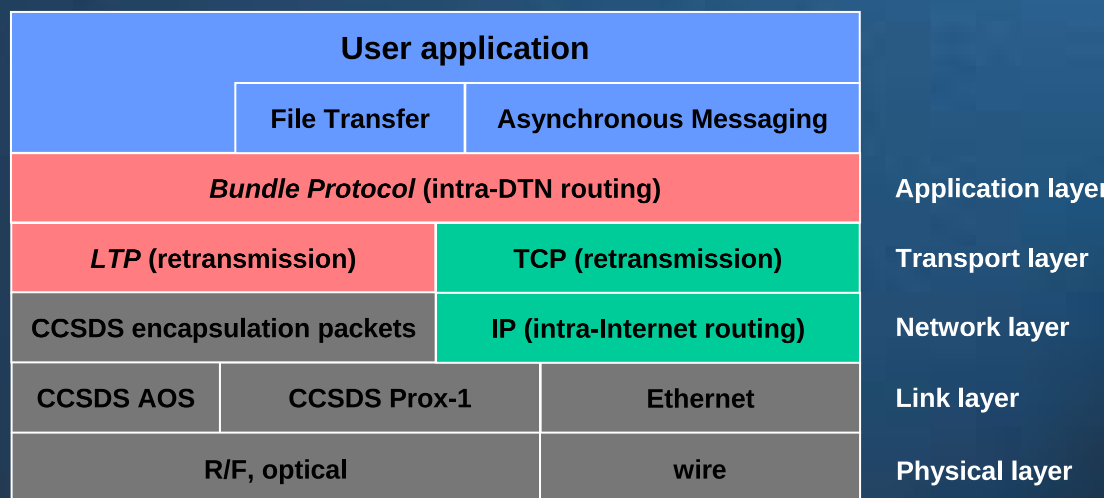

An End-to-end Path

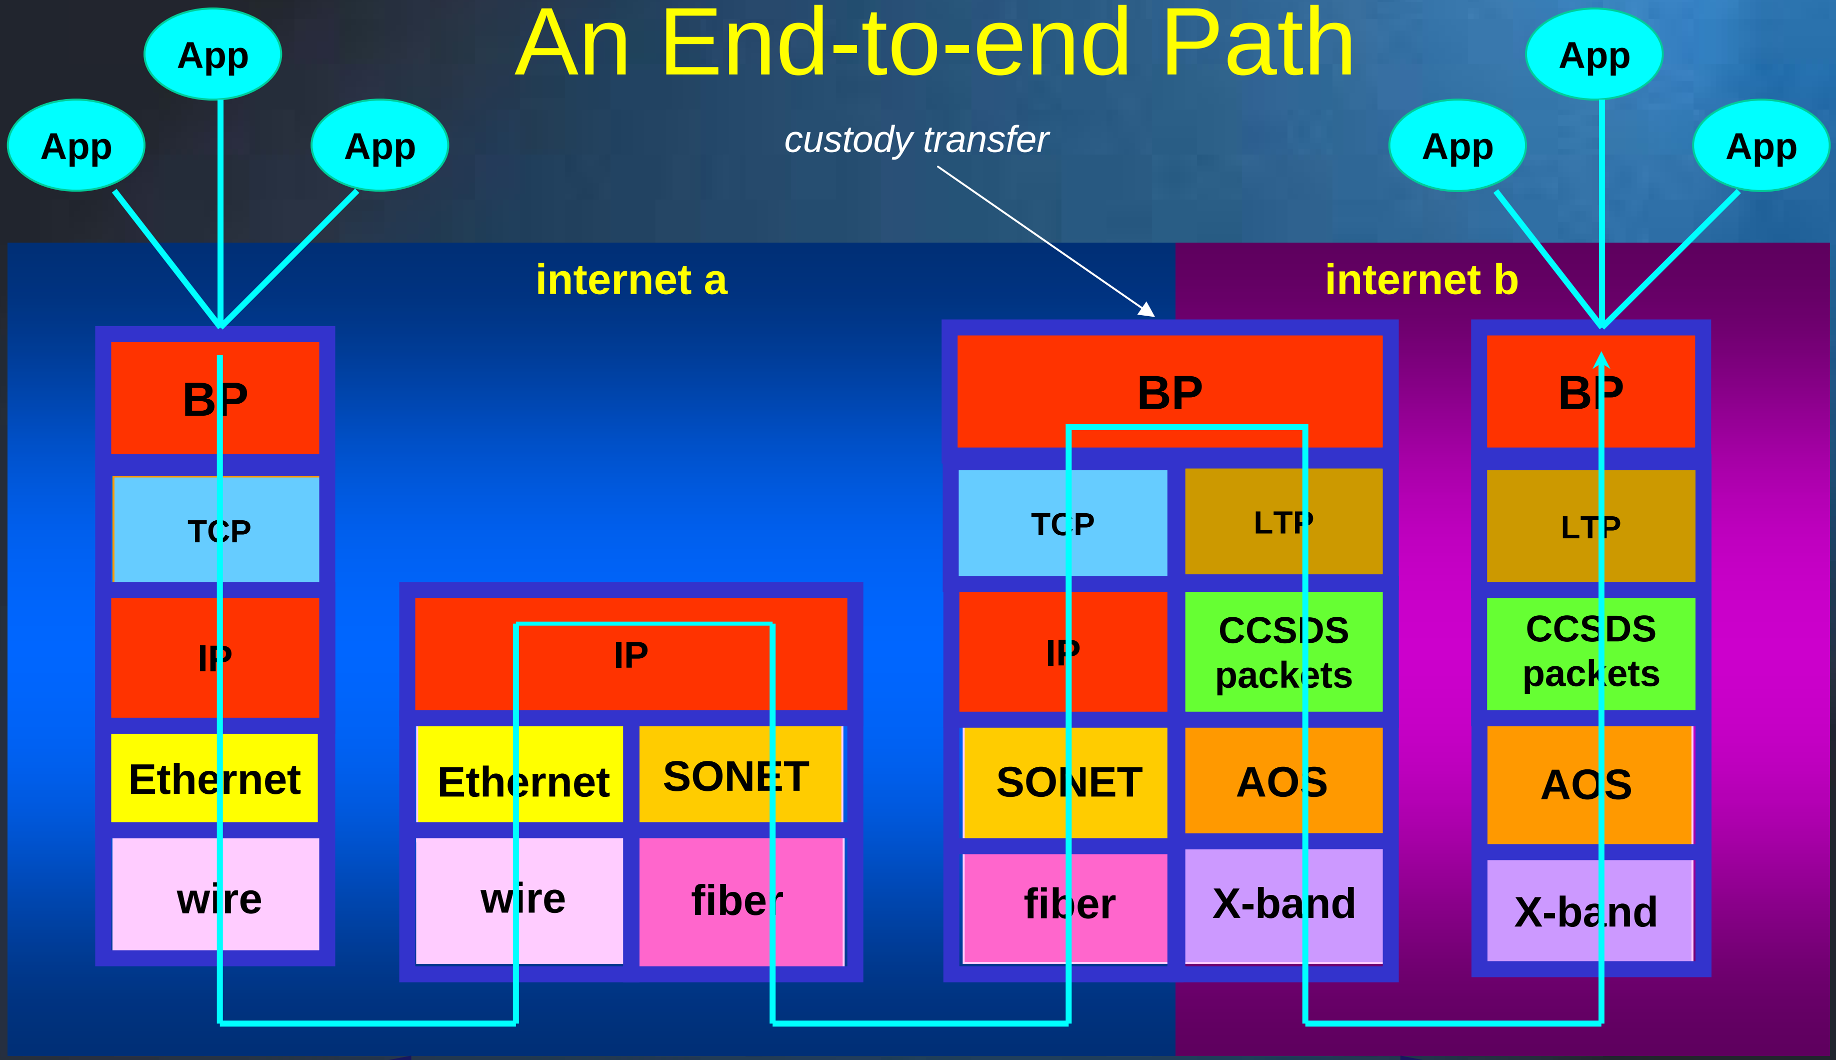

Network of internets spanning dissimilar environments

DTN Operations In Deep Space

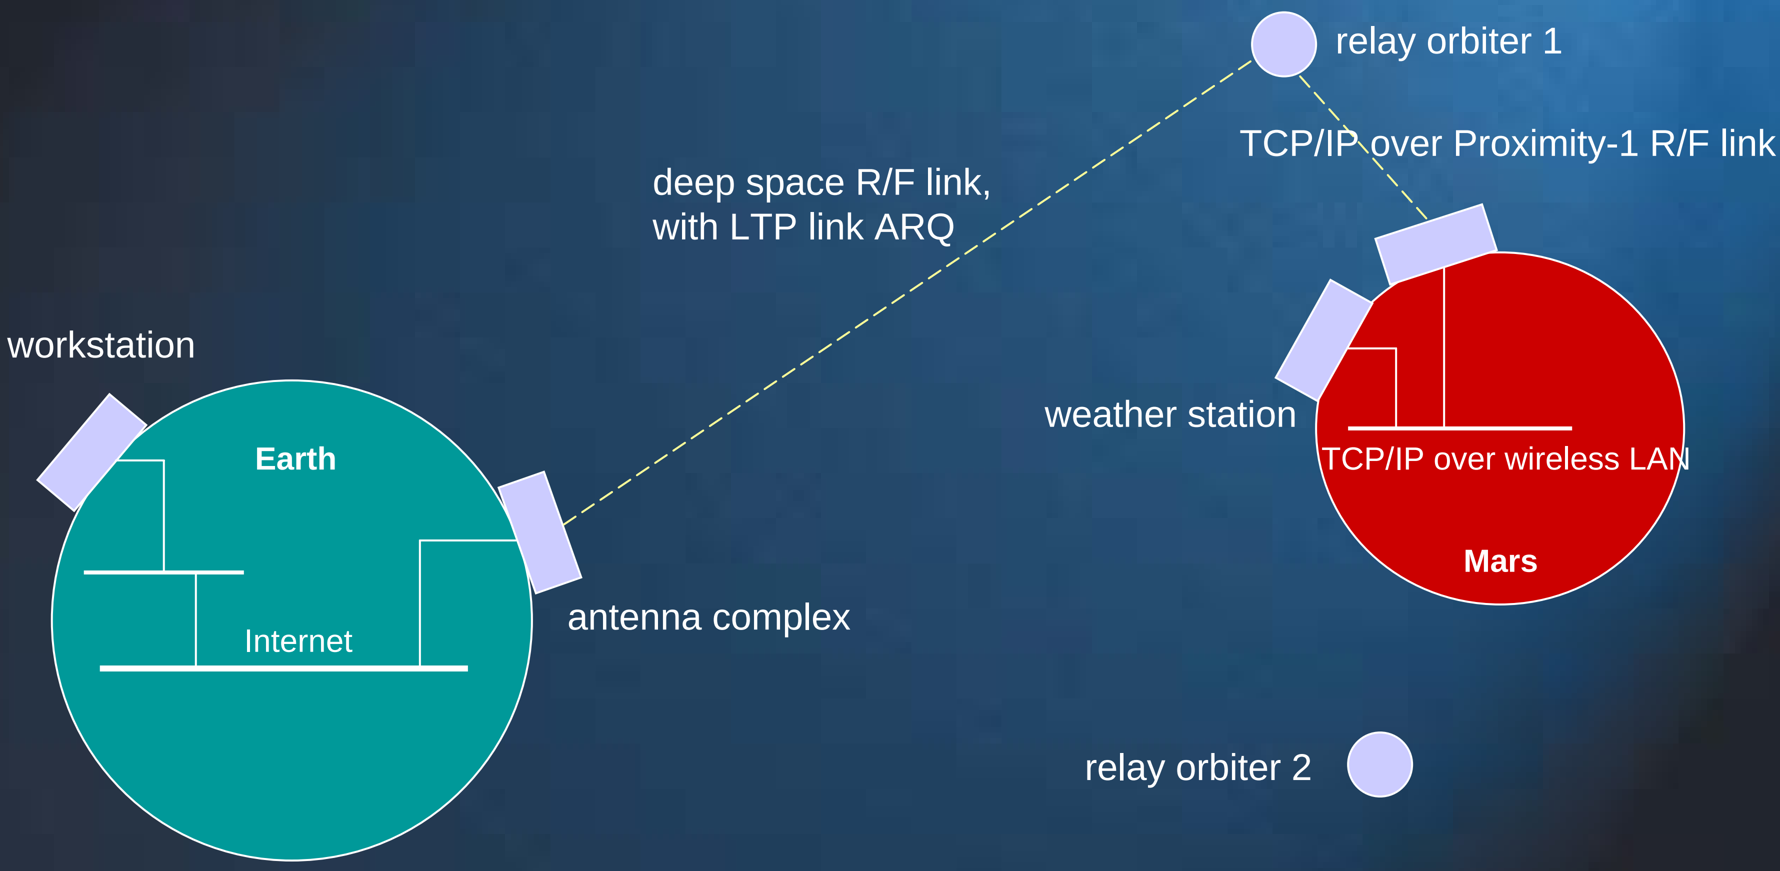

DTN Current Status

- Specifications and documentation
  - Internet Draft for the DTN architecture
  - Advanced Internet Drafts for both the BP and LTP protocol specifications
    - Plan to submit these as Experimental RFCs within IETF in 2006
- Implementations
  - BP implementations
    - DTN2: open source reference implementation (Intel, UC Berkeley)
    - ION: designed for space flight (JPL)
  - LTP implementations
    - Reference implementation in Java (Ohio University)
    - C implementation designed for space flight (JHU/APL)

Other Applications of DTN

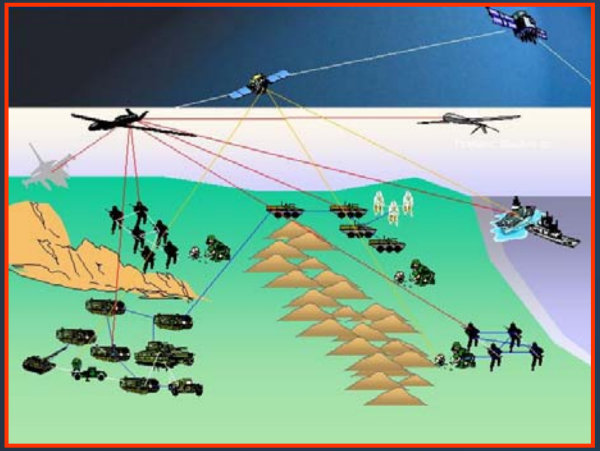

- Serving “challenged” regions of Internet topology
  - Mobile nodes: frequent disconnection, rapidly changing topology
  - Geographically remote nodes: relatively stable topology but infrequent connection
  - Undersea communication: long signal propagation delays, frequent disconnection
- Mobile tactical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

Remaining Problems

- 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 and ICMP source quench are end-to-end, may not reduce data injection rate at source until congestion collapse has already occurred.

Deep Space Experience to Date

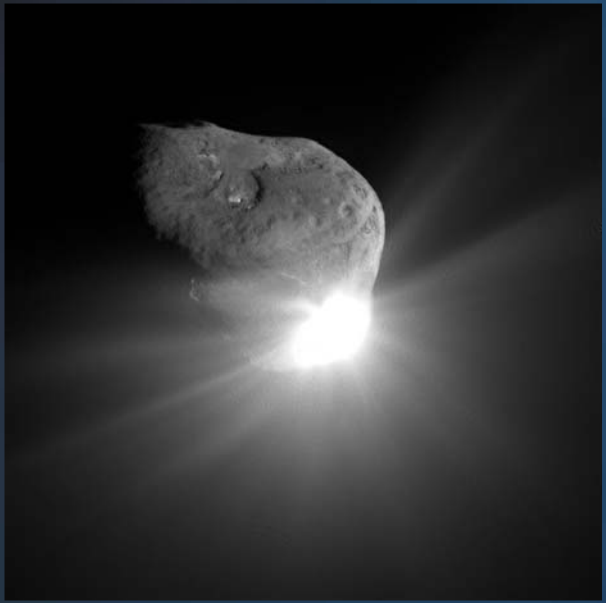

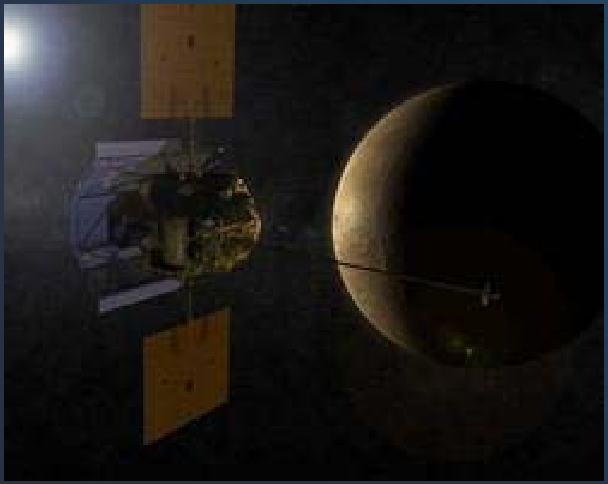

- CCSDS File Delivery Protocol is a precursor and functional prototype of DTN architecture.
  - Deferred, link-state-sensitive transmission
  - Delay-tolerant retransmission
  - Designed for relay over multiple space link protocols
- CFDP for Deep Impact:
  - 821 files uplinked
  - 133,000 files downlinked
- CFDP for MESSENGER:
  - about 5000 files downlinked
  - still in cruise to Mercury


What We’re Aiming For…