Interplanetary Overlay Network (ION): What's New

Scott Burleigh
IPN Group
18 May 2015

Provided through the courtesy of the Jet Propulsion Laboratory, California Institute of Technology.

Delay-Tolerant Networking

<u>The terrestrial Internet</u>

Continuous end-to-end connectivity
Very brief round-trip times

<u>The Solar System Internet</u>

Intermittent point-to-point connectivity
Very long round-trip times

Constraints on DTN in Space

	Terrestrial DTN	**DTN for Space Flight**
Links	Ethernet or WiFi (omni) Fast, cheap, symmetrical	Directed R/F, highly attenuated Relatively slow, very expensive, asymmetrical Must use reception/transmission contacts efficiently.
CPU, memory	Commodity generic chips Fast, cheap	Limited-production radiation-hardened chips Relatively slow, very expensive Must use processing resources efficiently.
Resource management	Reboots are easy. Dynamic management of memory is routine.	Hands-on repair is impossible; must minimize risk. Dynamic memory management is unpredictable. Fixed memory allocation is provided at startup.
Operating System	Commercial O/S with memory protection; tasks run in user space.	Real-time O/S, normally no memory protection – all tasks run in the same space as the kernel. Must be RTOS-compatible.

ION's Design

- The design of ION addresses these constraints.
  - Built-in private dynamic management of memory allocated at startup.
  - High-speed shared direct access to built-in object database.
  - System-wide transaction mechanism, for safety:
    - Ensures mutual exclusion, preventing lockouts and race conditions.
    - Enables reversal of all database updates made within the current transaction in case of software failure.
  - Compressed bundle headers, for transmission economy.
  - Zero-copy objects, for processing and storage economy.
  - Written in C, for processing economy and small footprint.
    - About 60,000 physical lines of code.
    - About 35,000 logical lines of code (omitting comments and whitespace).
  - Portable among POSIX operating systems, including RTOS.
    - Currently available for Linux, Solaris, OS/X, FreeBSD, VxWorks, RTEMS.
    - Also Windows and Bionic.

Design Overview

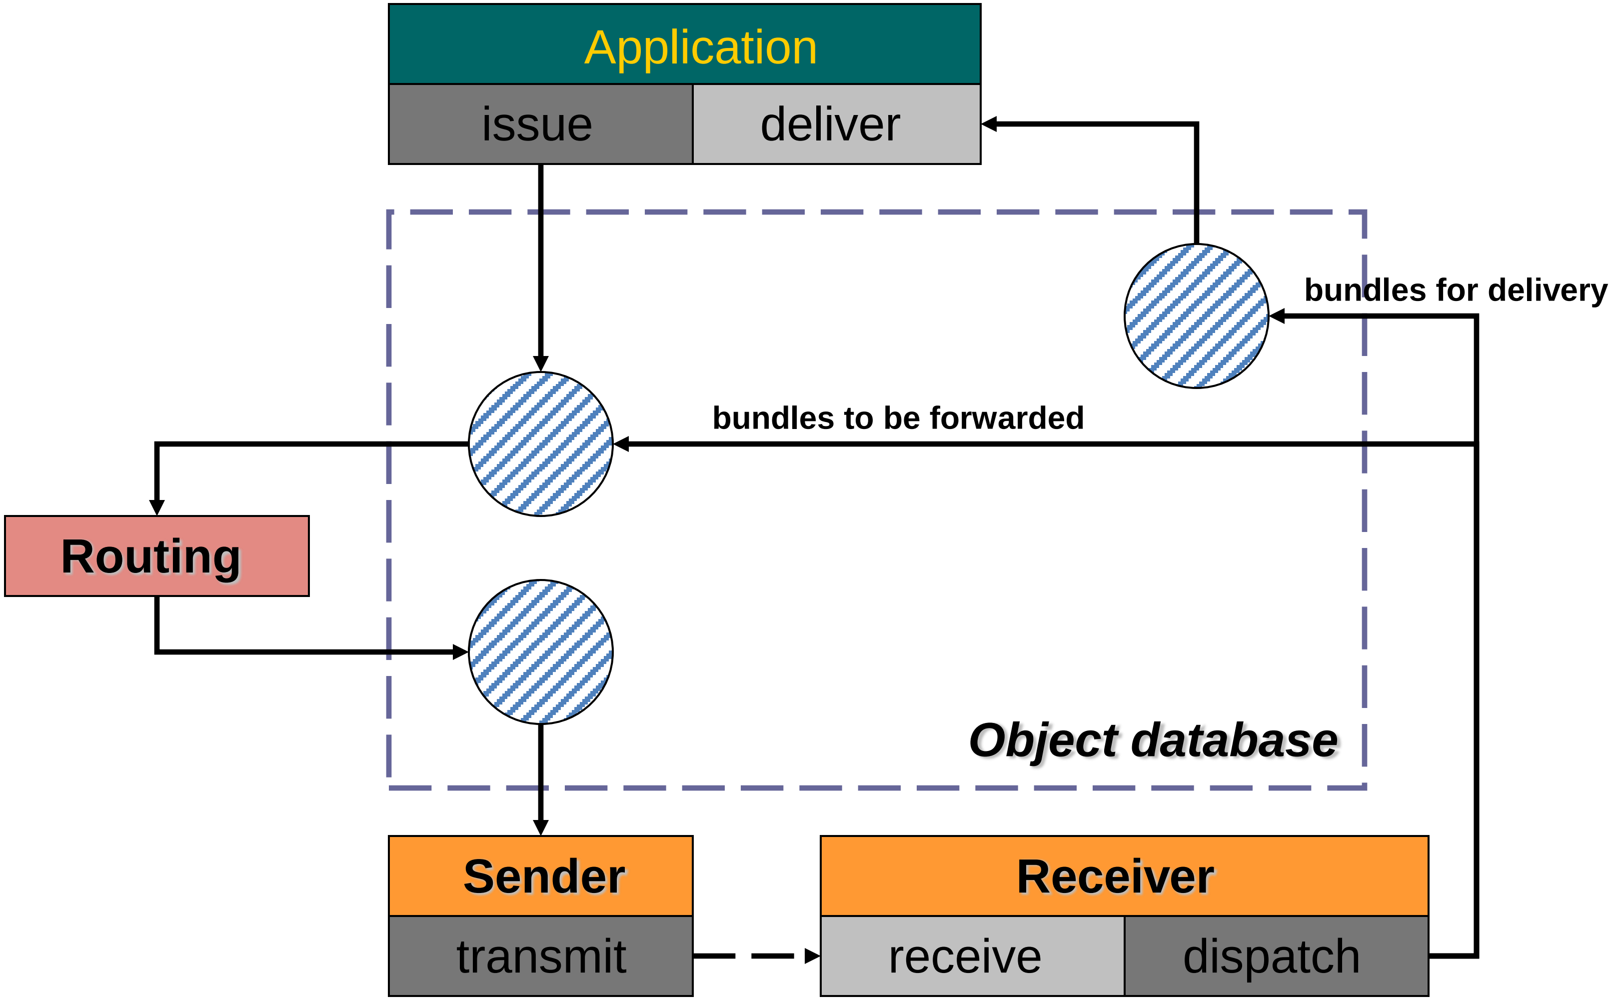

ION Design Principles

- Use shared memory.
  - Often there's no protected memory, so we have no option.
  - But this can be turned to advantage: shared memory is a highly efficient, very robust way to pass data between flight software tasks.

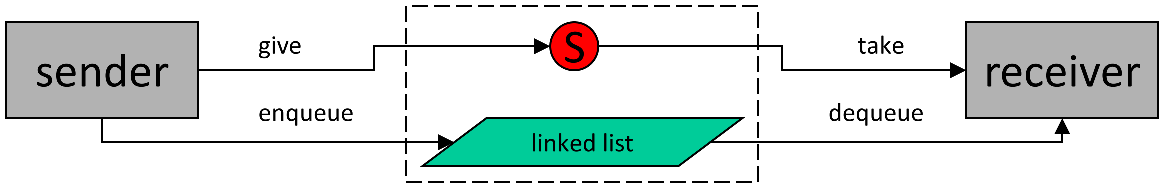

- Zero-copy procedures: leverage shared memory to minimize processing overhead.
  - Encapsulation in layers of protocol overhead (headers and trailers) can be done by reference rather than by copy.
  - The same data object can be shared by multiple tasks, provided reference counting prevents premature deletion.
- Portability: this is an unfamiliar programming model, so we must make it easy to develop in an environment with good programming support (e.g., Linux) and then deploy – without change – in the target RTOS environment.

Deep Impact Network Experiment (DINET)

- DINET: an experimental validation of "ION" (Interplanetary Overlay Network), JPL's implementation of the Delay-Tolerant Networking protocols. ION was uploaded to the backup flight computer of the EPOXI spacecraft on 18 October 2008 and was operated continuously from that date until 13 November 2008.
- Spacecraft functioned as a DTN router in an 11-node network (all other nodes on Earth).
- One-way signal propagation delay was initially 81 seconds (24 million km), dropped to 49 seconds by the end of the four-week exercise. Transmission to spacecraft at 250 bytes/second. Transmission from spacecraft at either 110 or 20000 bytes/second.
- Moved 292 images (about 14.5 MB) through the network, no data loss or corruption.
- First deep-space node on the Interplanetary Internet:
  - Automatic, contact-sensitive relay operations (store-and-forward Bundle Protocol)
  - Automatic rate control
  - Delay-tolerant retransmission (Licklider Transmission Protocol)
  - Prioritization of merged traffic flows
- Longest digital communication network link ever.
- First use of dynamic routing over deep space links.
- Demonstrated that fully automatic operation of a delay-tolerant network over deep space links is feasible.

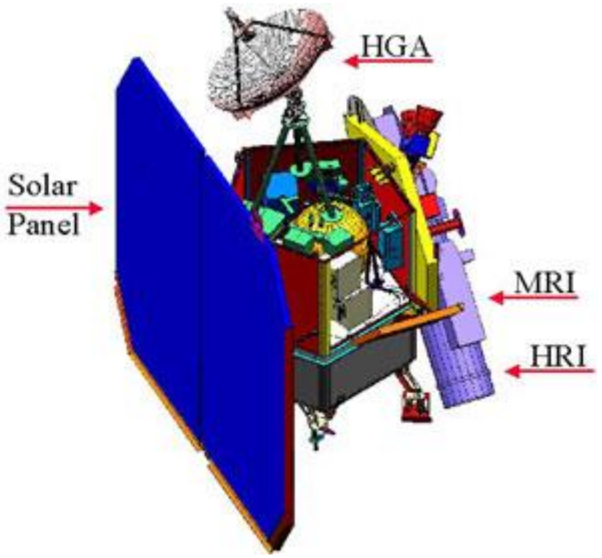

EPOXI (formerly Deep Impact) spacecraft

A Brief History of ION

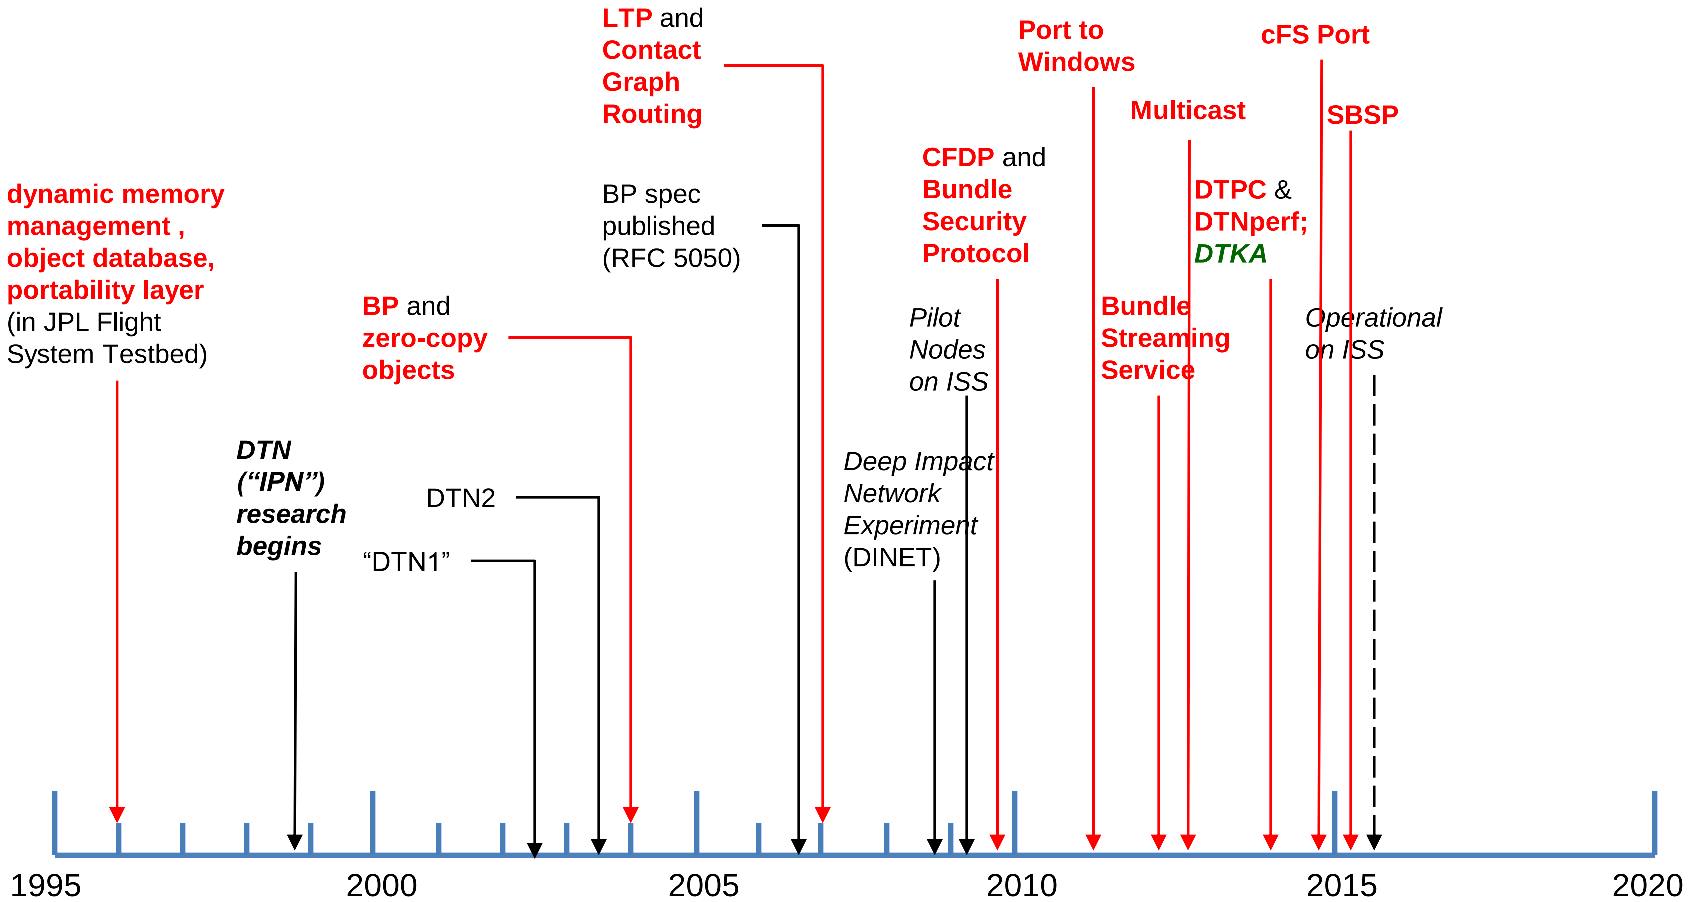

Current Status (1 of 2)

- ION 3.3.0 released 4 March 2015, focus on security.
  - Integrates Streamlined Bundle Security Protocol into ION.
    - As proposed in Internet Draft; not yet adopted as a standard by IETF.
  - Integrates Bundle-in-Bundle Encapsulation into ION.
  - Implements dynamic computation of reporting limits in LTP, to enable transmission of large blocks despite high rates of packet loss.
  - Improves support for operations on International Space Station (ISS).
    - Flow control on Simple TCP convergence-layer adapter, providing ability to handle inbound and outbound traffic levels that always exceed available bandwidth.

Current Status (2)

- ION 3.3.1 released 30 April 2015.
  - Additional features and bug fixes needed for ISS:
    - Windows: semaphore management, bping port.
    - Dynamic re-binding of TCP connections.
  - Dynamic computation of LTP retransmission limits.
  - Bug fixes resulting from CCSDS BP interoperability testing, between ION and DTN2.
  - Updates to the Tutorial.

Focus for ION 3.4.0: Terrestrial DTN

- Opportunistic routing:
  - Contact discovery protocol.
  - Contact plan exchange protocol.
  - Probabilistic contact inference.
  - Extension of contact graph routing to operate on probabilistic contacts.
  - Ground operations are integral to flight missions.
- Possibly get Delay-Tolerant Key Administration released as open source and integrated into ION.
- Targeted for early July 2015.

Focus for ION 3.5.0: Scalability

- Federated node registration infrastructure:
  - Node registration servers (maybe based on LDAP):
    - Protocol for propagating node registry information.
  - Node auto-configuration:
    - Initially use MAC address as own node number.
    - Send registration information to registration server.
    - Get registration information, as needed, from server.
  - Multicast is key – notify newly created nodes of multicast group membership.
  - Extend routing to concept of “regions”.
- Targeted for early October 2015.

On the Horizon

- Develop implementations of revised BP and BSP specifications as they emerge from IETF DTN Working Group deliberations.
- Develop a spanning-tree maintenance protocol, in support of BP multicast.
- Leverage probabilistic contact plans:
  - Provide congestion forecasting in opportunistic networks (support for congestion control).
  - Use contact plan to estimate bundle delivery latency, use as value of bundle’s time-to-live (more support for congestion control).

Somewhat Over the Horizon

- Registration infrastructure configures new nodes?
  - Maybe initial node number and region number are just temporary, and permanent node number and region number are automatically assigned by registration server.
  - DHCP-ish.
- Develop Delay-Tolerant Information-Centric Networking, like a “self-forming Akamai.”

Backup slides

Flight Environment Constraints (1 of 3)

- Link constraints: wireless links enabling interplanetary network communication are generally slow and are usually asymmetric.
  - Limited electrical power, relatively small antennae.
    - So signals are weak. This limits transmission from the spacecraft to rates on the order of .25 Mbps to 6 Mbps.
    - Additionally, reception sensitivity is limited. Rates of transmission to the spacecraft are typically even lower, on the order of 1 or 2 Kbps.
  - So the cost per octet of data is on the links is high, and the links are heavily subscribed.
  - Economical use of reception and transmission opportunities is important.

Flight Environment Constraints (2 of 3)

- Processor constraints:
  - Limited electrical power, limited mass allowance.
  - Intense radiation environment, mandating radiation-hardening, which is time-consuming and expensive.
  - Relatively small market, limiting incentive to do radiation-hardening engineering for the latest advances in processor technology.
  - So flight processors are always slower than engineering workstations.
  - So the cost per processing cycle is high and the processors are heavily subscribed.
  - Economical use of processing resources is important.

Flight Environment Constraints (3 of 3)

- Hands-on repair is impossible, so reliability is key.
  - Predictability enhances reliability, so flight software usually must meet hard real-time deadlines. So real-time operating systems (RTOS) are used: all software runs in “kernel” (rather than “user”) mode, no memory protection.
  - Dynamic allocation of system memory is difficult to predict, so it is typically prohibited except in certain well-understood spacecraft states, e.g., start-up.

Software Elements (1 of 3)

- Interplanetary Communication Infrastructure (ICI)
  - “platform” library insulates ION software elements from the differences among operating systems.
  - Personal Space Management (PSM) enables flexible, dynamic private management of a fixed block of pre-allocated system memory.
  - Memory Manager system enables coexistence of multiple memory management instances (e.g., multiple PSM-managed partitions).
  - Lyst and SmList systems standardize management of linked lists in private and shared memory. SmRbt system manages red-black trees in shared memory for high-speed indexing and retrieval.
  - Simple Data Recorder (SDR) enables flexible, dynamic private management of a fixed block of non-volatile storage, such as battery-backed memory or a pre-allocated file in a flash file system.
  - Zero-Copy Objects (ZCO) system enables protocol encapsulation by reference rather than by copy and provides a reference counting system to enable safe concurrent access to a single non-volatile storage object by multiple tasks.

Software Elements (2 of 3)

- Licklider Transmission Protocol
  - Full implementation of the LTP spec as developed by the DTN Research Group. (RFC 5326)
  - Additional features:
    - Aggregation of multiple service data units into a single block, to minimize the volume of acknowledgment traffic over highly asymmetric links.
    - Implements delay-tolerant, non-conversational flow control based on limiting block size and the number of transmission sessions that can be in progress concurrently.

Software Elements (3 of 3)

- Bundle Protocol
  - Full implementation of the BP spec as developed by the DTN Research Group. (RFC 5050)
  - Includes support for:
    - Prioritization of data flows
    - Bundle reassembly from fragments
    - Flexible status reporting
    - Custody transfer
  - Additional features:
    - Rate control provides support for congestion forecasting and avoidance.
    - Bundle headers are compressed, to reduced protocol overhead and improve link utilization.
  - Also includes an implementation of Contact Graph Routing, a system for dynamic routing over interplanetary links.

Core Components of ION

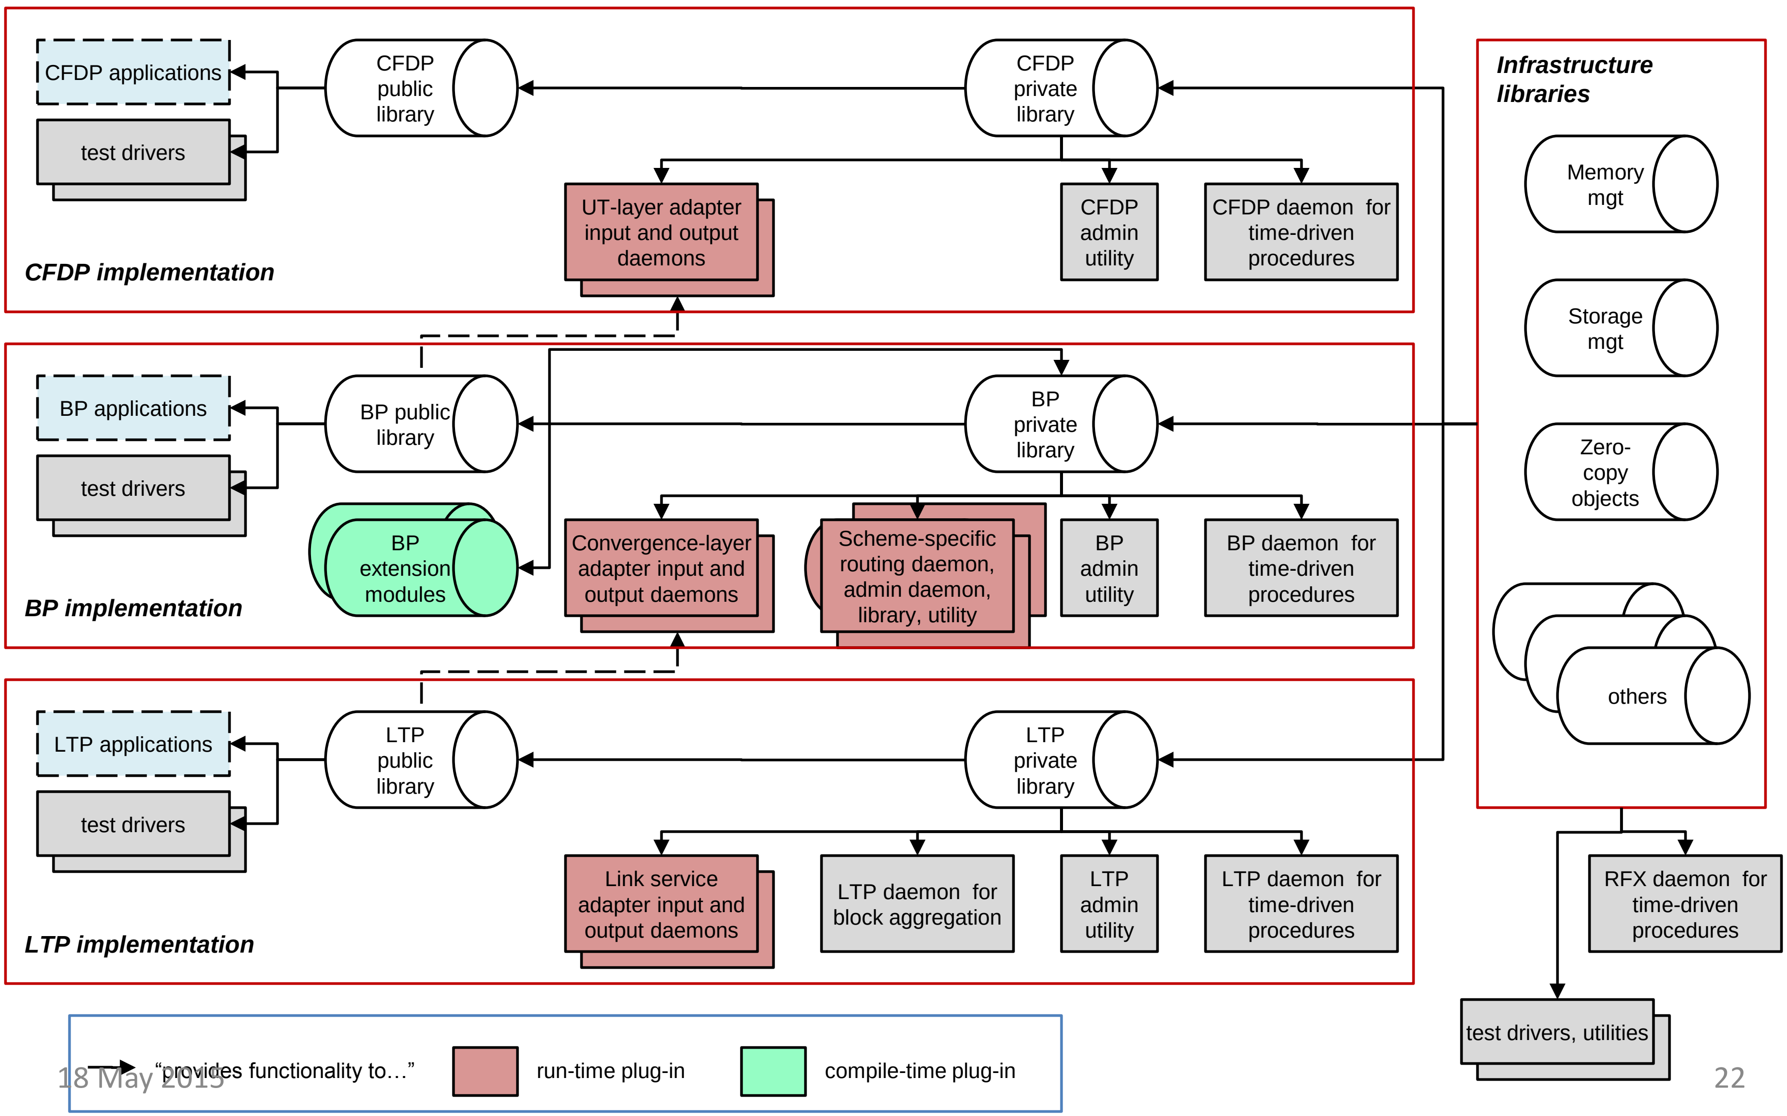

Compressed Bundle Header Encoding

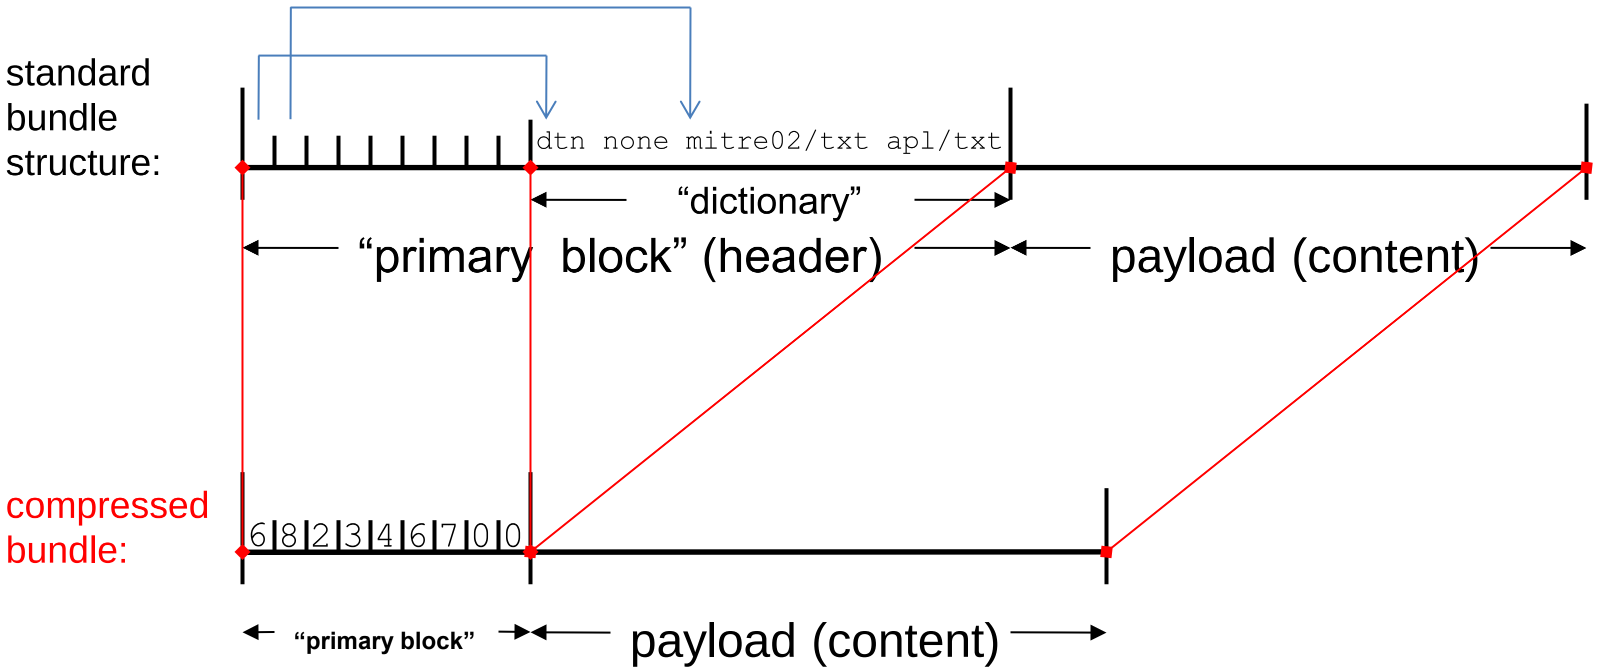

Managed Aggregation

BP custody transfer (or CFDP)

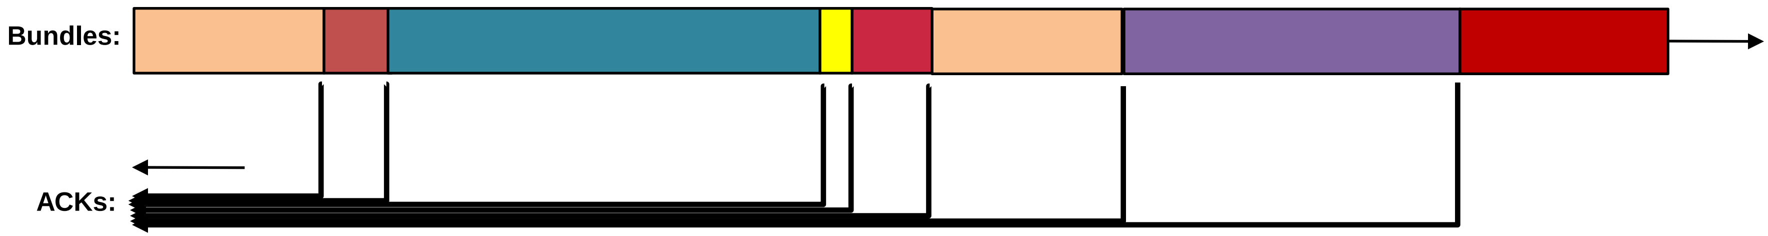

LTP in ION

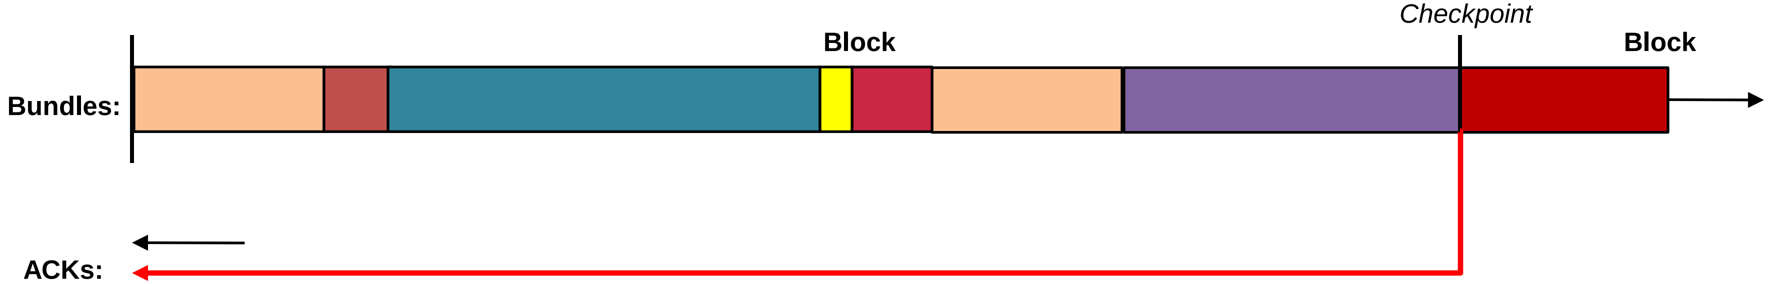

Block size is configurable, so ACK rate can be tuned to the return data rate.

Private Dynamic Management of Fixed, Pre-allocated Memory

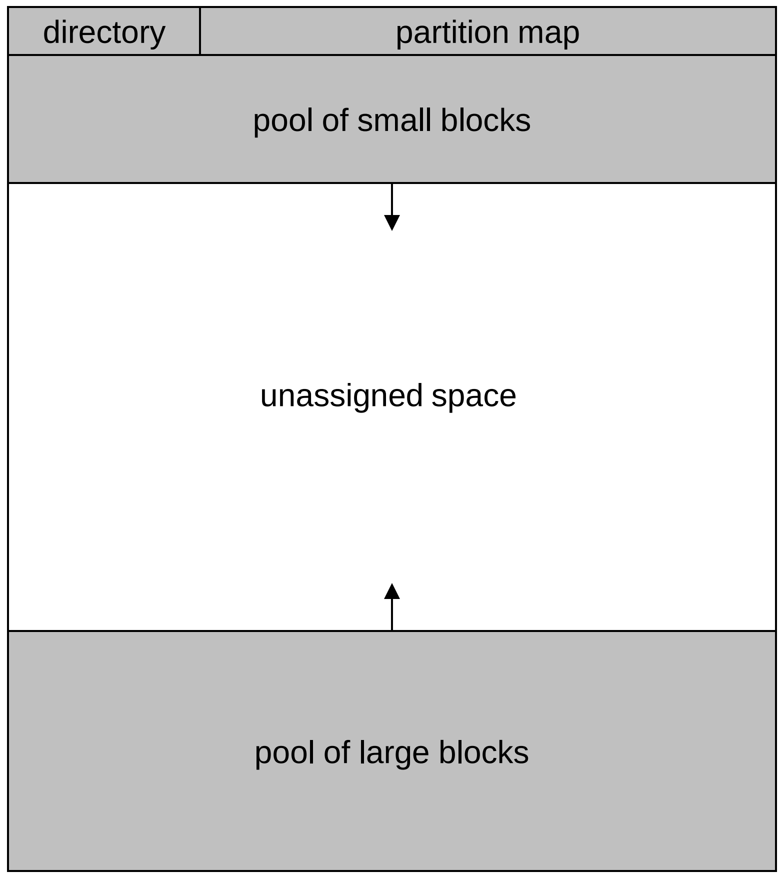

Note: this system has been in continuous use for JPL projects since 2004, on the EO-1 spacecraft for the Autonomous Science Experiment and on autonomous sea surface and undersea vehicles for the Navy.

Non-Volatile Data Store

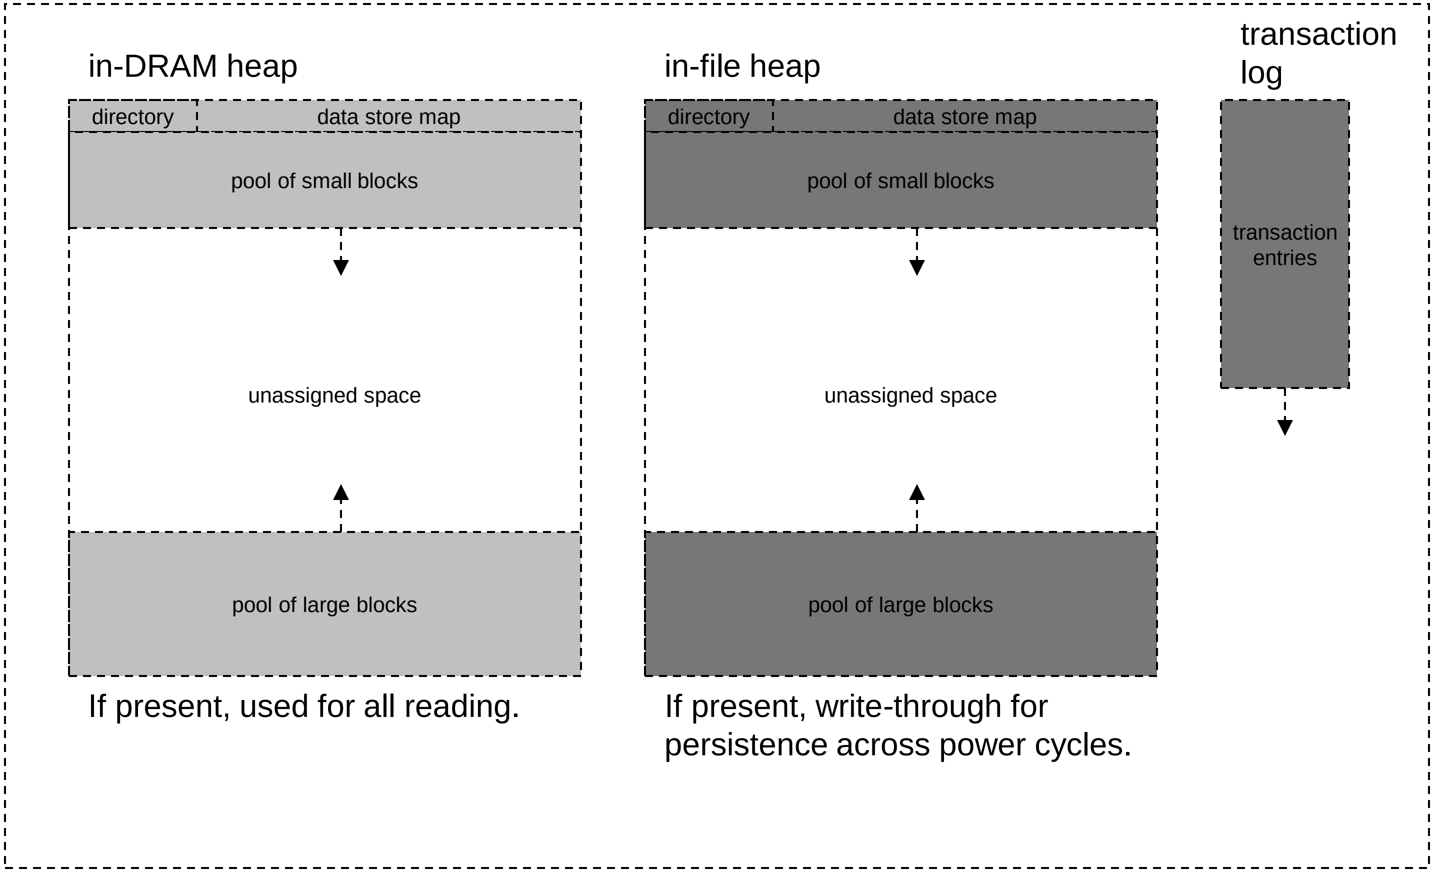

Memory Management Blocks

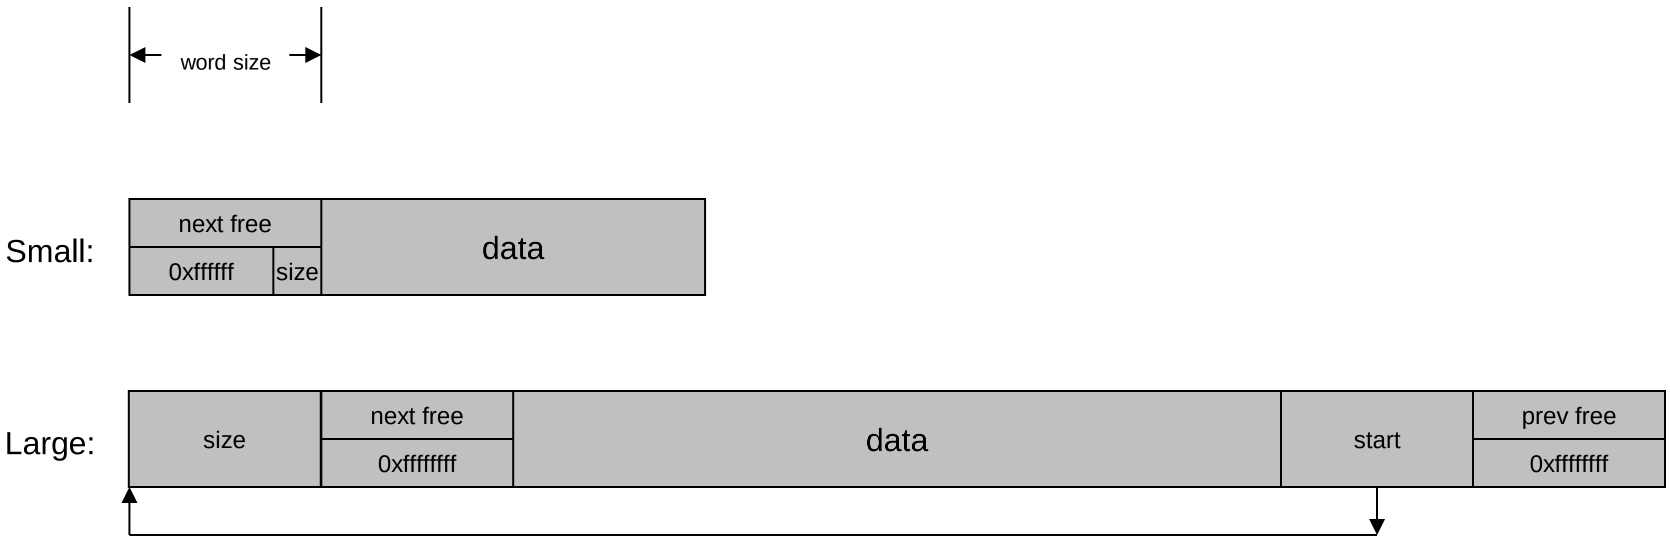

Trailing overhead of large block enables a newly freed block to be merged with the adjacent free block(s), if any, to minimize fragmentation.