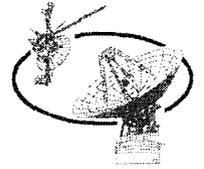


CCSDS File Delivery Protocol in Delay-Tolerant Networking

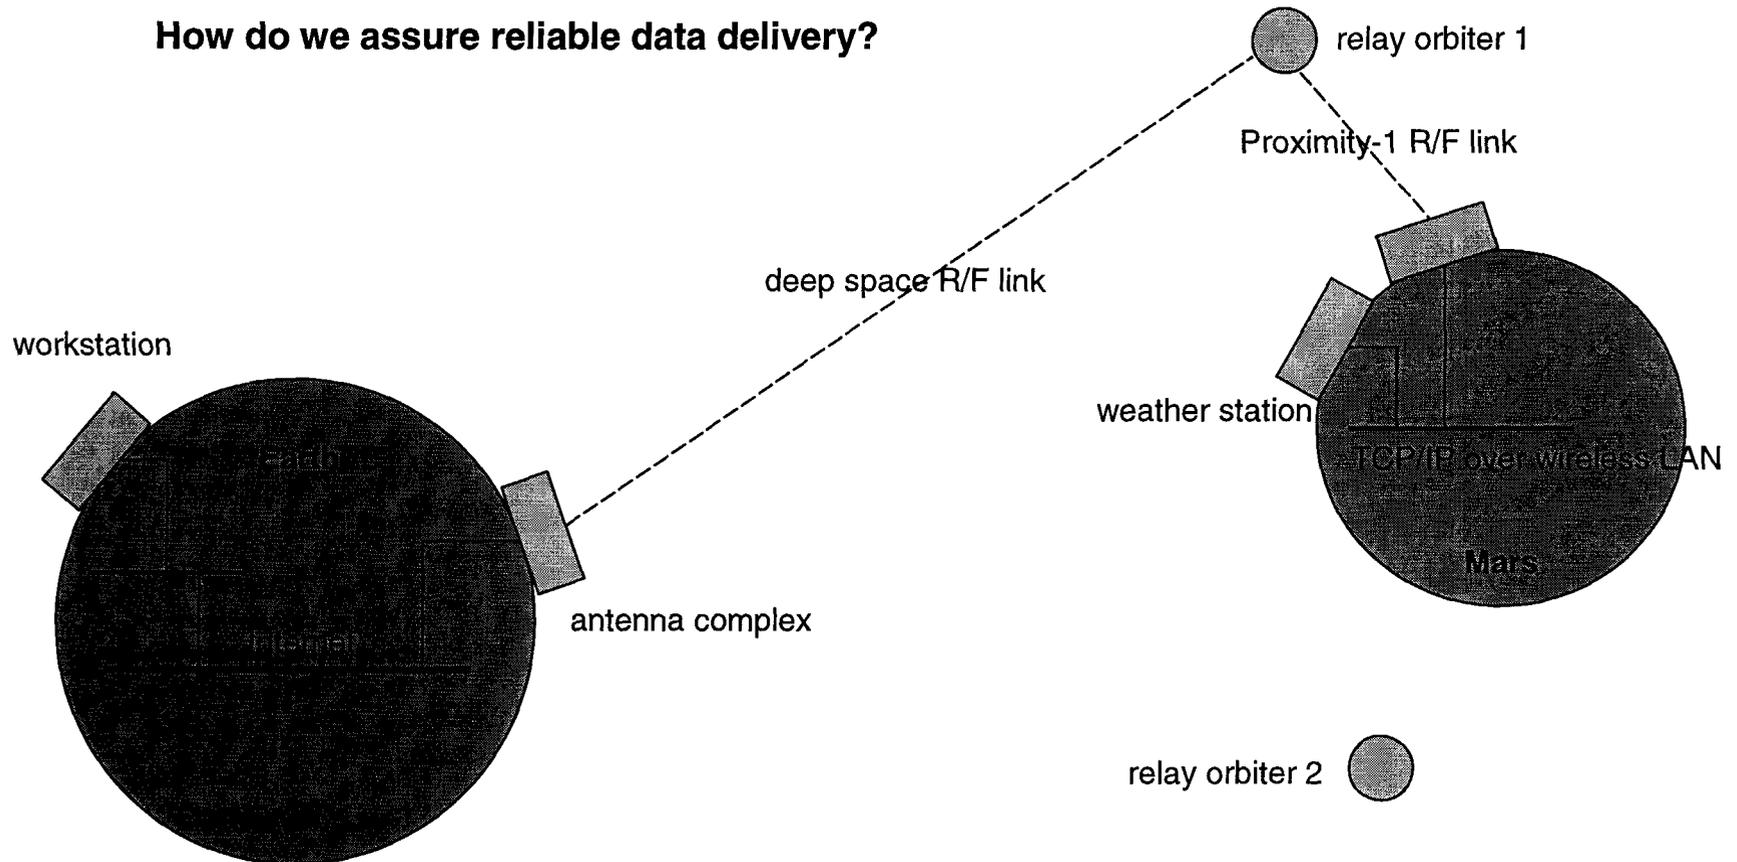
Scott Burleigh
NASA Jet Propulsion Laboratory
4 March 2003

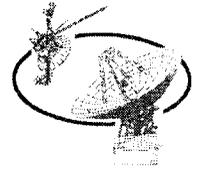


A Mars Operations Scenario



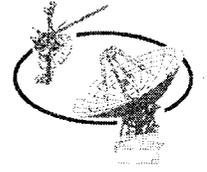
How do we assure reliable data delivery?



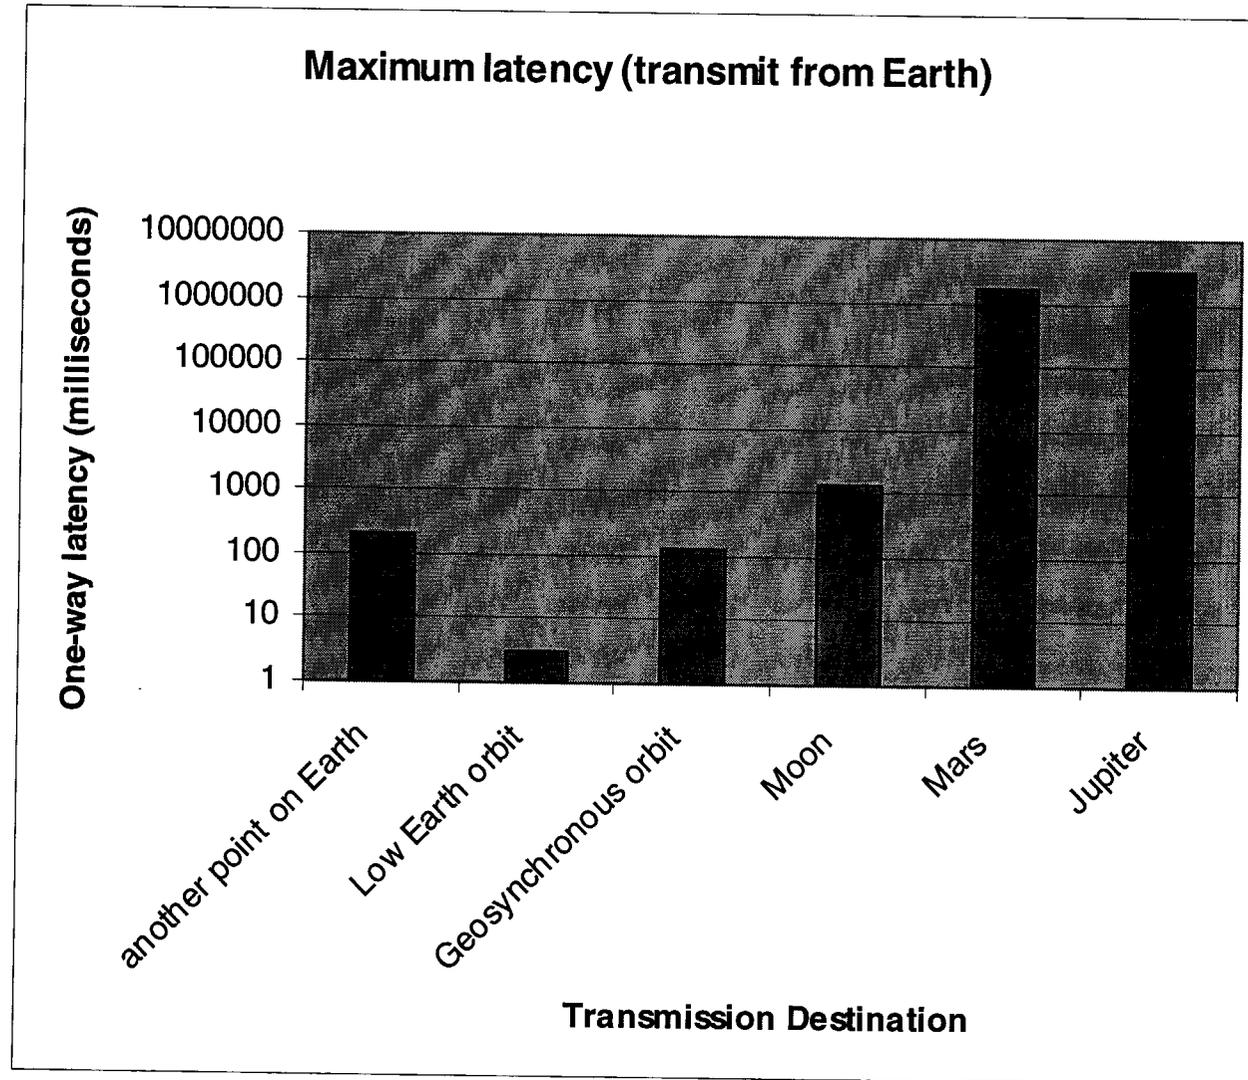


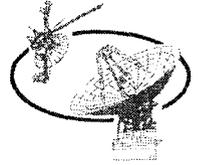
Deep Space Communications

- Spacecraft have limited power and antenna size, so data rates are low and are asymmetrical.
- Links are noisy due to solar wind, etc.
- But the central problem is extremely long round-trip communication times:
 - Intermittent connectivity. For example, the DSN may “track” some spacecraft for only 2 hours per day – or 8 hours once per week.
 - Very long distances, fixed speed of light, so signal propagation delay is on the order of minutes or hours rather than milliseconds.
- Reliable transmission of any single byte can take an arbitrarily long time:
 - Transmission can be lost due to corruption, N times.
 - NAK can be lost due to corruption, N times.
 - Connectivity can be lost between time of transmission and time of reception, so transmission of NAK (or of data) in response can be delayed by hours or days.



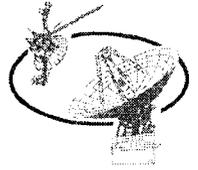
Latency in Interplanetary Communications





Implications for End-to-end Protocol

- Connection establishment could take days.
 - So protocol must not rely on dynamic connection establishment or negotiation.
- In-order stream delivery could suffer arbitrarily long periods of paralysis, waiting for byte N to be received before delivering byte N+1.
 - So out-of-order delivery is needed.
 - So protocol must support multiple transmissions in flight concurrently.
 - So data must be structured in messages (transmission blocks) for accountability and concurrent retransmission; not in streams.
- But any single message transmission 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 to minimize risk of massive transmission failure.

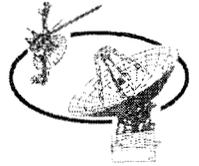


More Implications

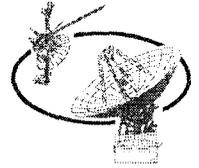
- Continuous end-to-end transmission through relay elements may be impossible, due to time-disjoint episodes of connectivity.
 - So relays can't just route packets; they must store them, and then forward them when opportunities arise.
- End-to-end retransmission would reserve resources (retransmission buffer) at originator for entire duration of the transaction – possibly days or weeks.
 - So retransmission should be point-to-point rather than end-to-end. “Custody transfer.”
 - So ARQ is needed at every relay point. The links themselves need to be reliable.



CFDP is...

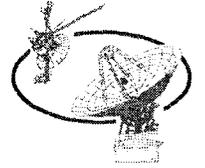


- An international standard for automatic, reliable transfer of files between spacecraft and ground (in both directions), built on top of the familiar CCSDS protocols.
- Monolithic – it includes both application and transport (reliability) functionality in a single protocol – yet part of a layered architecture that makes it ...
- A foundation technology for additional deep space communication capabilities built into user applications:
 - File retrieval (“Get”).
 - Proxy file transfer.
 - Remote directory query.
 - Remote transaction suspension, resumption, status query.
 - Others as needed.

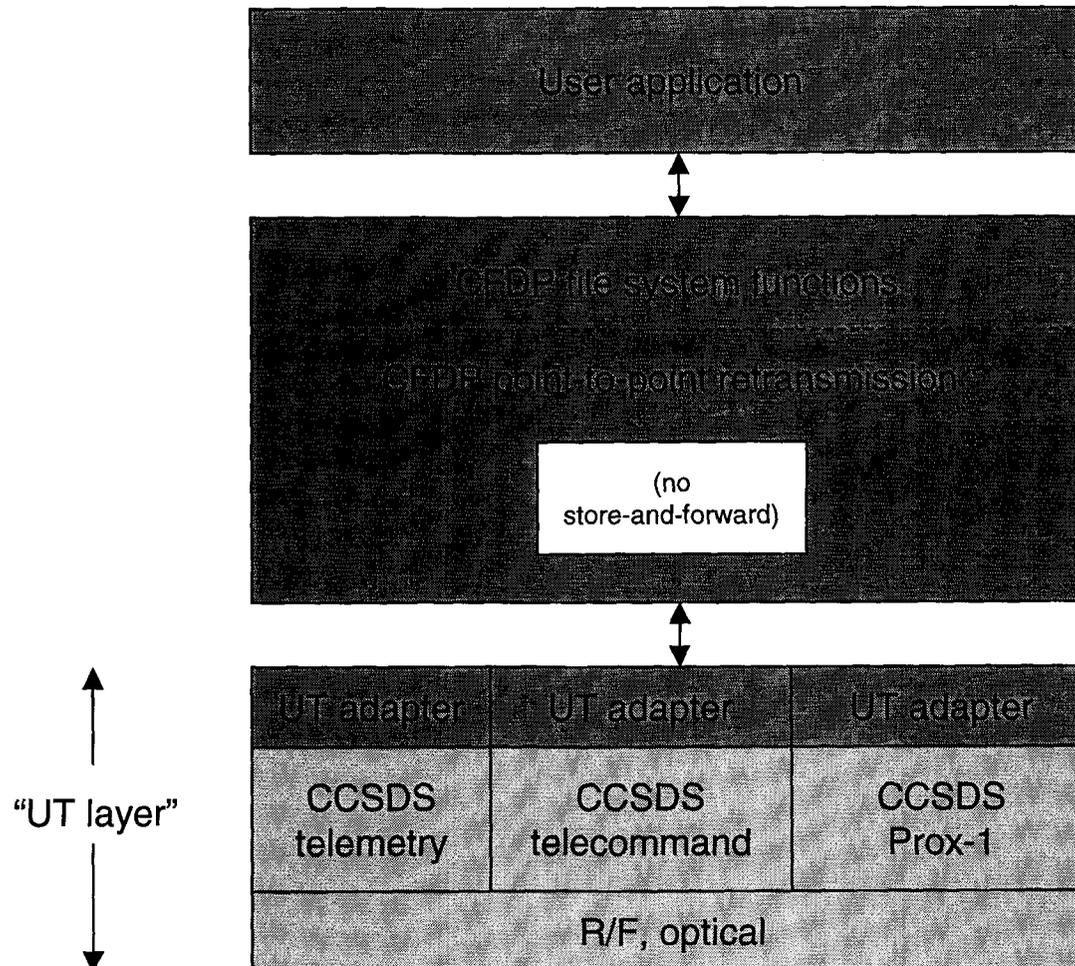


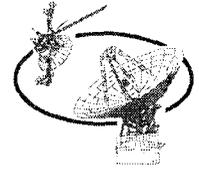
CFDP Core Operations

- Premise: entities can communicate directly (R/F or optical).
 - Mutual line-of-sight visibility.
 - Compatible operating schedules: entity A can point at entity B and transmit at a time when entity B can point at entity A and receive.
 - Adequate links: the levels of transmitter power and receiver power combine to produce a data rate greater than zero.
- Transfer of files between two entities over interplanetary distances.
- Metadata can be associated with each transaction.
 - Small messages to user applications.
 - Remote file system operations commands: cp, mv, rm, mkdir, etc.
- Concurrent transfer transactions, multiple retransmission buffers; incremental (possibly out of order) delivery.
- Deferred transmission: outbound data are stored pending transmission opportunity.
- Delivery is reliable – protocol automatically retransmits lost or corrupted data.



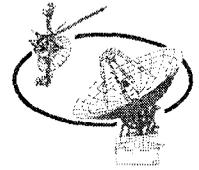
Core Architecture





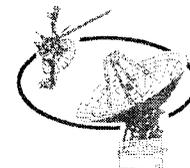
CFDP Benefits

- Efficient operation over simplex, half duplex, and duplex links, interplanetary distances.
 - No dynamic connection protocol.
- Effectiveness over highly unbalanced link bandwidths.
- Effectiveness over a range of mission profiles, from low Earth orbit to deep space.
- Application can request a file transmission at any time, without knowledge of when the communication link will be available.
- Portions of file can be made available to the user as soon as they arrive.
- Minimized link traffic, due to aggregated selective negative acknowledgments.
- Transfers can span ground station contacts (time disjoint connectivity).
- Transfers can span multiple ground stations that are acting as frame or packet forwarders.

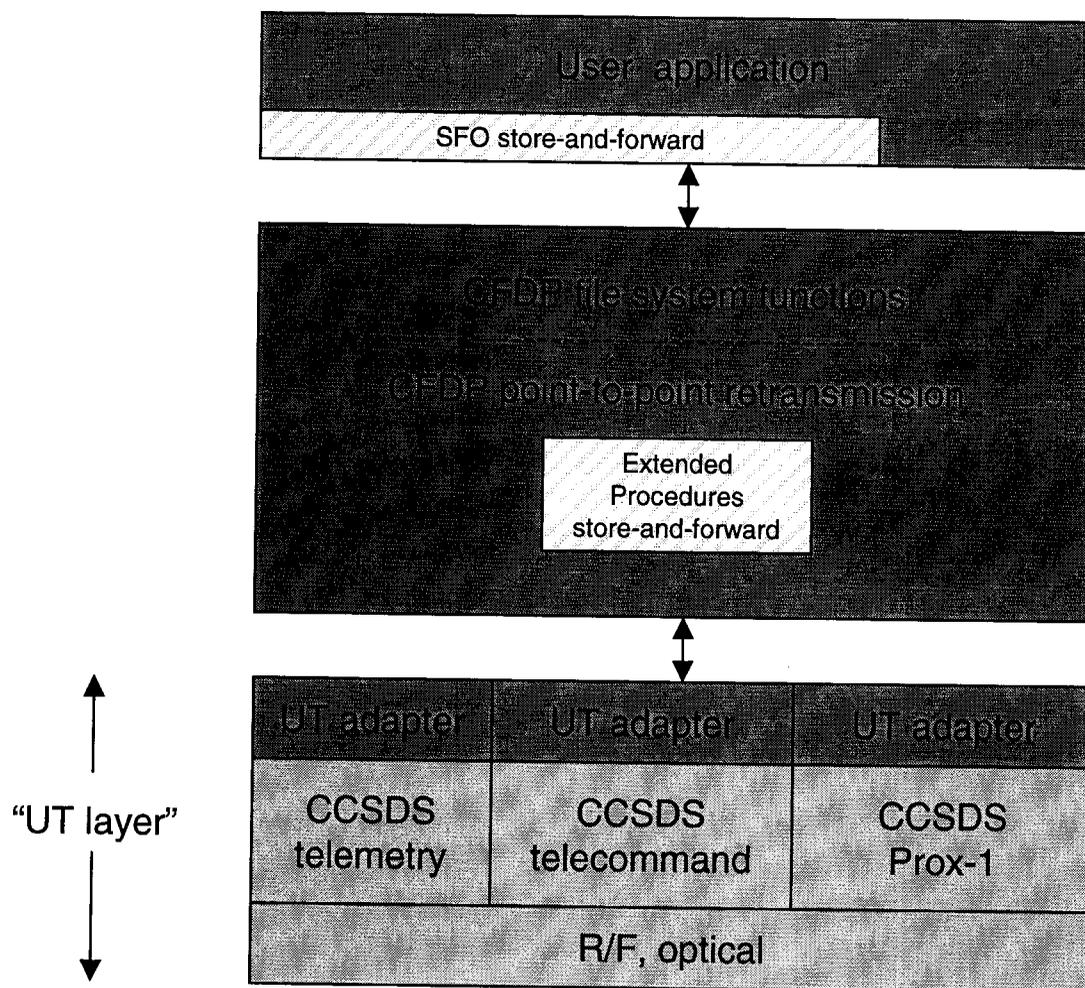


CFDP Extended Operations

- Premise: entities cannot communicate directly.
 - No mutual visibility: intervening planetary mass, intervening Sun.
 - Incompatible operating schedules.
 - Insufficient signal power between sender and receiver.
- So CFDP supports indirect communication, via “relay” or “waypoint” entities:
 - Simple, static routing algorithms.
 - Storage of data at waypoints, pending outbound connection.
 - Retransmission loop is closed between waypoints, not between endpoints.
- Implementation options:
 - Extended procedures
 - Additional functionality built into CFDP itself.
 - Implemented by ESA; implementation at JPL planned for FY03.
 - Store-and-forward Overlay
 - CFDP is left unchanged.
 - Additional functionality built into standard user application layer.
 - Implemented by JPL, may be implemented by ESA as well.



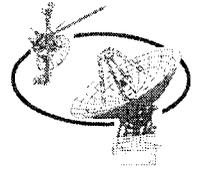
Extended Architecture





Benefits of Extended CFDP

- All the benefits of Core CFDP.
- Can additionally support more complex operating scenarios, such as the “Mars Operations” example given above.
 - Reliable end-to-end transmission of a file over multiple links,
 - one or more of which may be deep space links,
 - no two of which need be concurrently active.
 - Notice of delivery to ultimate destination (end-to-end acknowledgment).

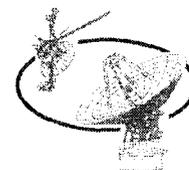


Extended CFDP vs. Internet

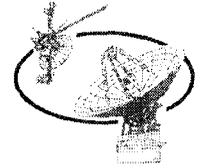
- However, extended CFDP is not an ideal general solution for communication in the “Mars Operations” scenario:
 - Designed and optimized for file transfer. Little support for other styles of communication – messaging, streaming video or audio, database query.
 - Each protocol engine can only be used by a single user application.
 - No support for dynamic routing, so network size is limited.
 - No end-to-end security.
- The Internet protocol suite (TCP/IP) has none of these flaws, but it's not the ideal general solution either:
 - Retransmission is only end-to-end.
 - Routing protocols rely on continuous connectivity to all points in the network: sustained link outage – even if temporary and scheduled – is assumed to mean the entities on the far side of that link are unreachable.



Communication Environments



	<u>Workstation to Ground Station</u>	<u>Ground Station to Relay Satellite</u>	<u>Relay Satellite to Weather Station</u>
<i>signal propagation latency</i>	milliseconds	minutes	milliseconds
<i>data rate</i>	10-100 Mbps	8-256 Kbps	8-256 Kbps
<i>communication mode</i>	bidirectional	time-disjoint transmission and reception	bidirectional
<i>connection mode</i>	continuous	intermittent; scheduled	intermittent; opportunistic
<i>network access</i>	on-demand	managed	managed
<i>congestion potential</i>	high	none	low



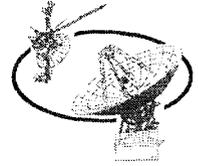
The General Problem

- Question: how do we best enable communication between applications running at two arbitrary locations in the Interplanetary Internet?
- Constraint: the answer cannot assume that any of the following conditions hold in the end-to-end path.
 - continuous connectivity
 - low or constant transmission latency
 - low error rate
 - low congestion
 - high transmission rate
 - symmetrical data rates
 - common name or address expression syntax or semantics
 - data arrival in transmission order



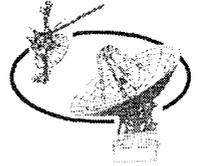
A General Solution

- Our answer: insert an overlay protocol at a new *delay-tolerant networking* (DTN) layer of the protocol stack, between the application and transport layers.
- The overlay protocol, called *bundling*, unifies multiple Internets (and other types of networks designed for specific environments) in the same way that the Internet protocol IPv4 unifies multiple LANs (and other types of subnets).



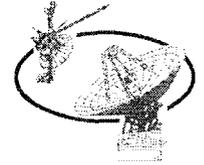
DTN Architectural Principles

- Postal model of communications
 - Asynchronous transmission, minimal conversational interchange. Avoid query/response model, negotiation.
 - Transmitted data are self-contained, self-identifying units of work: data bundled together with requisite metadata – *bundles*. (Hence “bundling”.) Somewhat like e-mail messages with attachments.
- Tiered functionality
 - Protocols that operate well within each environment already exist.
 - DTN *nodes* are entities that communicate using the Bundling protocol. A *region* is a set of nodes that are mutually reachable using a given environment-specific protocol family.
 - Rely on the capabilities of regional protocols as much as possible. Use Bundling (at the layer above) to bridge between regions.
- Terseness
 - Bandwidth is not universally cheap. Don’t waste it.



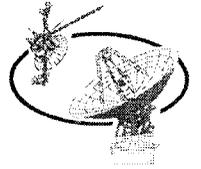
Tiered Forwarding

- Regional network-layer protocols, such as IP in Internet-like regions, forward bundles among the DTN nodes within each regional network.
- At the boundary between two regional networks, a DTN *gateway* node operates like an IP router but at the DTN layer:
 - receives bundles within the “upstream” region via its regional protocol;
 - forwards them within the “downstream” region via its regional protocol.
- Connectivity is not guaranteed to be continuous, so a node may be unable to forward bundles immediately.
 - Must instead store them until a forwarding opportunity arises – *deferred transmission*.
 - May need to use a non-volatile medium (instead of DRAM) for storage, to conserve DRAM and increase protocol robustness.



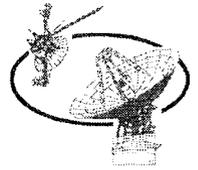
Tiered Naming and Addressing

- In order to be forwarded, and perhaps answered, bundles must be tagged with source and destination node identifiers.
 - Each node is within some region, so node identifier must include regional node identifier for regional forwarding.
 - Source and destination may be in different regions, so node identifier must include region ID for DTN forwarding.
- A DTN node identifier is a *tuple* comprising both region identifier and also a regional node identifier (such as DNS name) that can be mapped to a regional address:
 - {region ID, regional node identifier}
- Tuple need not contain regional address, such as IP address.
 - Regional node identifiers are administrative, not topological.
 - They are *late bound* (mapped to addresses at moment of delivery), to insulate remote bundle sources from local topological change.



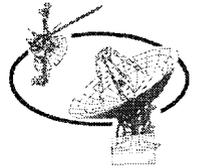
Tiered Routing

- Multiple possible routes (both in space and time) to the destination.
 - Route to next hop within the same region – if not point-to-point – is computed by region-specific protocol, such as IP within the Internet.
 - End-to-end routes are computed by Bundling.
- DTN routing decisions must be sensitive to anticipated transmission opportunities (*contacts*):
 - continuous connectivity in some regions
 - scheduled contacts, e.g. in deep space
 - Schedules might be loaded via management interface or routing protocol.
 - discovered (opportunistic) contacts in low-latency regions
 - predicted contacts, based on analysis of contact history



Tiered ARQ

- Transport-layer ARQ provided by regional protocol assures reliable delivery from one node to the next.
 - In Internet-like regions, TCP.
 - In the interplanetary link (IPN backbone) region, an adaptation of CFDP's core retransmission procedures.
- DTN-layer ARQ supplements transport-layer ARQ as needed (e.g., in case a relay node crashes):
 - Optional transfer of custody performed within Bundling.
 - Not all forwarding nodes need be custodians.
 - Optional end-to-end reception report, retransmission.



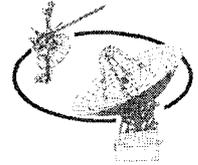
Tiered Security

- To minimize unauthorized consumption of DTN resources (bandwidth, memory, processing cycles), relay nodes will not forward bundles received from non-authenticated nodes – that is, DTN nodes are *mutually suspicious*.
 - Sender's encrypted credentials are included in bundles.
 - Keys for encryption and decryption are distributed in certificates.
 - Certificates may be pre-placed, may be distributed periodically in encrypted traffic, or may accompany the bundles to be forwarded.
- The same certificate distribution mechanisms may be made available to applications, for end-to-end security of application data.



Tiered Congestion Control

- Regional congestion avoidance and control:
 - In low-latency regions, will typically be accomplished by feedback loops built into the protocols.
 - E.g., within the Internet, TCP includes congestion avoidance algorithms.
 - In high-latency regions, will typically be accomplished by management – that is, competition for link access is resolved by reservation rather than contention.
 - E.g., in deep space communications, congestion is avoided during operations planning rather than in real time.
- This may be sufficient. If not, a supplemental DTN-layer congestion control mechanism may be needed.



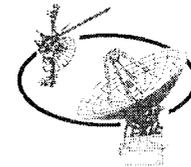
Resilient Delivery

- The ultimate destination of a bundle is a *service agent*, a running task or process or thread that provides or consumes an information service.
- Latency in a DTN end-to-end path may be so great that the destination service agent is not running at the moment a bundle destined for it arrives.
- So the final destination node may need to store a bundle until the service agent for which it's destined is running and able to receive it: *deferred delivery* in addition to deferred transmission.
- The final destination node may even need to start (or restart) the destination service agent itself so that a bundle can be delivered: *reanimation*.



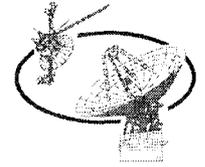
Postal Service Levels

- Three levels of delivery priority: low, standard, and high.
- Three postal service notifications, all of which can optionally be sent to a specified “reply-to” service agent rather than to the original sender:
 - Notice of initial transmission, i.e., notice of mailing.
 - Notice of delivery to the ultimate destination application , i.e., return receipt.
 - Report of route taken, i.e., delivery record.

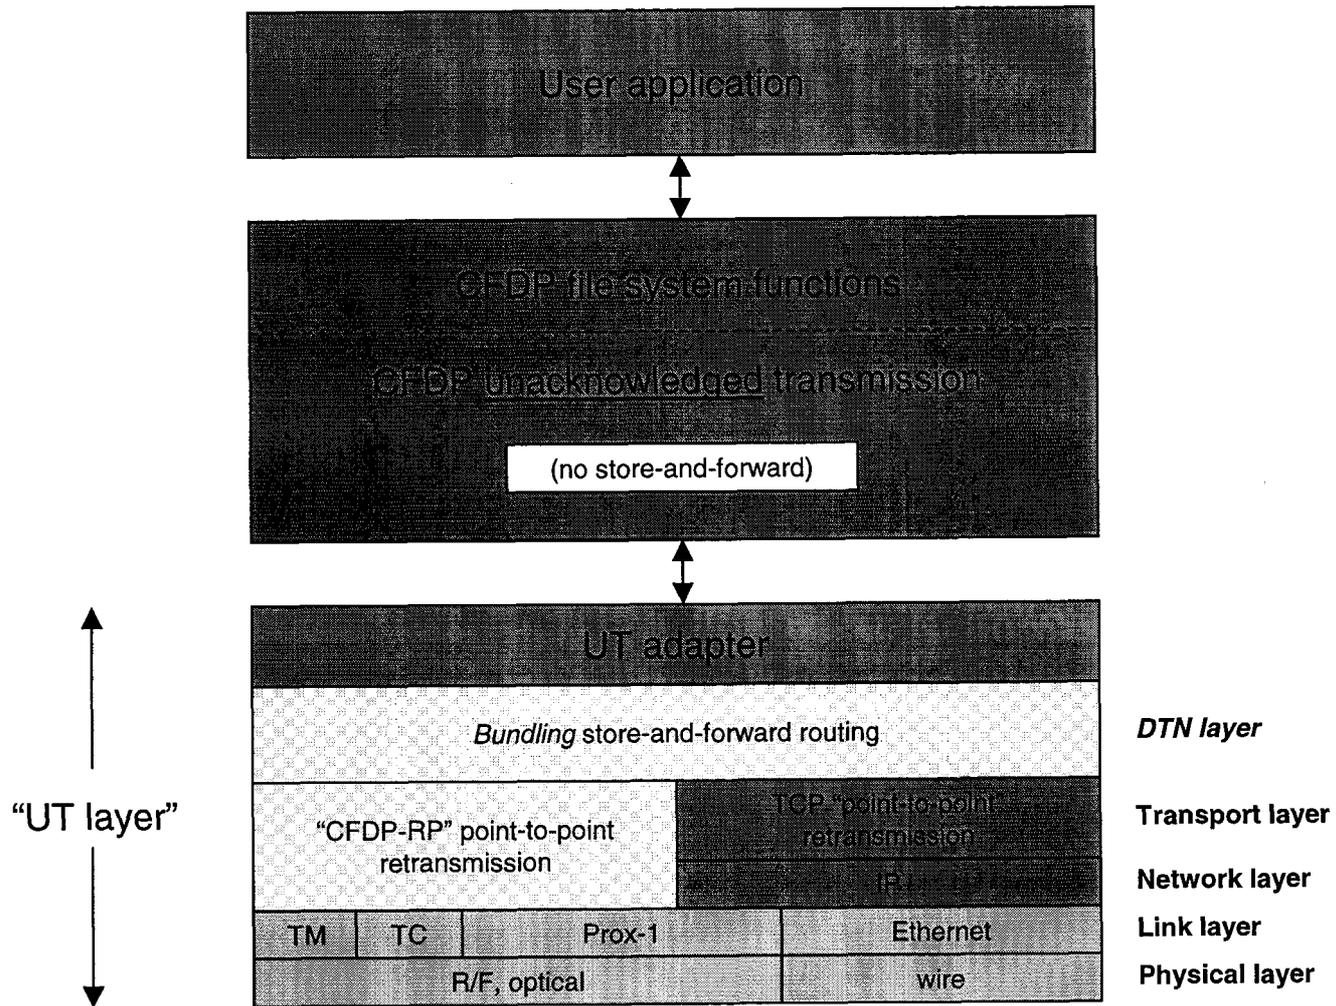


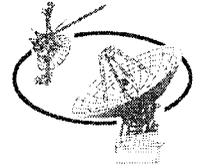
CFDP in the IPN

- A fair question: if Bundling does all this, why would anyone still need CFDP?
- The two-part answer:
 - The long-latency retransmission procedures that constitute core Acknowledged CFDP are still needed.
 - Extract them into a regional transport protocol (“CFDP-RP”) that assures reliable direct transmission between two adjacent nodes within the Deep Space region.
 - Use Bundling to bridge between that protocol and TCP/IP.
 - The application-layer elements of CFDP are not duplicated by anything in Bundling. CFDP remains a critically valuable application protocol for deep space mission operations.



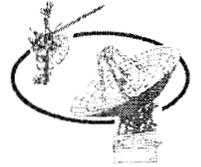
CFDP/DTN Architecture





CFDP: “Killer App” for the IPN

- Transmission of files: reads from one filestore, writes to another.
- File segmentation for effective transmission: awareness of record structures, e.g., files of CCSDS packets.
- Optional incremental delivery of portions of files as they arrive.
- Automatic file checksum verification, other fault handlers.
- Extensible file metadata mechanism: messages to user.
- Remote filestore management commands synchronized with successful file delivery.
- Standardized user operations:
 - File transmission by proxy, including “get”.
 - Remote directory listings.
- In short, CFDP provides all the services that are needed for file-based mission operation, while Bundling and regional protocols provide the services that CFDP itself needs.



CFDP: Looking Ahead

- A stable, internationally accepted mission operations standard.
- Supports reduced-cost mission operations based on reliable file transfer and remote file system management.
- A highly capable deep-space communications system in its own right.
- Can become even more capable and powerful as the DTN-based Interplanetary Internet grows in scope and complexity.