Proposed Evolution of CCSDS Protocols - Strategy and Technology

Richard D. Carper[1], Scott Burleigh[2], Adrian J. Hooke[2], J. Leigh Torgerson[2]

[1]*Consultant, Space Data Systems, Corvallis, Oregon.*

[2]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California.

Background

For more than two decades the international Consultative Committee for Space Data Systems (CCSDS) has developed and promulgated space link standards for the use of the international space-faring community. Nearly every member of that community is a member of the CCSDS. The organization has been sufficiently effective in its work that at the present time over 250 spacecraft have used or are using these standards. Using platforms range from high altitude balloons through all sizes and types of earth orbiters, including the International Space Station, through interplanetary spacecraft and planetary landers.

Currently some of the CCSDS member Agencies, most specifically the Jet Propulsion Laboratory (JPL) of the California Institute of Technology, are involved in the development of standards which allow the extension of the Internet into the space data domain. That domain consists of multiple regions which constitute a disconnected, variable-delay environment, and there is a growing need for a general, standard way to communicate end-to-end through this environment. The fundamental concept which JPL is using is that local in-situ short delay Internets, distributed across the Solar System on free flying spacecraft and on and around other planets, are interconnected via a long delay deep space backbone network. Just as the TCP/IP suite unites the Earth's Internet as a "network of networks", a new protocol suite called "Bundling" unites the InterPlaNetary Internet (IPN) into a "network of Internets".

JPL is testing the architectural concepts by beginning the detailed definition and prototyping of Bundling, and has charted a strategy whereby today's space systems can evolve smoothly into the Bundling era throughout the coming decade. It is especially noteworthy that the scope of the IPN has also been broadened so that it is a member of a new class of "Disruption Tolerant Networks" (DTNs).

Disruption Tolerant Networks

In many environments – such as interplanetary communications – a real time end-to-end path between two users may never exist as a connected entity. Instead, communications must occur by stitching together a series of time-disjoint individual hops. The time-disjoint character of point-to-point connections in a DTN may range from known, predicted losses of line-of-sight contact between orbiting vehicles and planetary rovers, through the unpredictable change of location of individual

mobile units, to pre-emption of resources in a battlefield situation or destruction of one or more intermediate nodes in the network. While deep space communications (with their enormous signal propagation delays) are obvious examples of DTNs, we expect that this mode of operation will become increasingly important to terrestrial communications as well. It seems evident that with the tremendous increase in the use of pocket size communicating devices of many different kinds and of similar vehicular devices, the networks which support them will inherently be DTNs. The result of this Internet evolution is that user devices are becoming increasingly mobile, and the end-to-end data path is traveling across a wider variety of environments. Mobility coupled with power, weight, volume constraints suggests that end-to-end communications dialogs across time disjoint connectivity are going to be highly likely, if not inevitable. The DTN protocol mechanisms could significantly enhance the Internet's ability to accommodate the extreme heterogeneity that is emerging, by bridging between islands ("regions") of homogeneity to support end-to-end delivery of information.

Therefore we believe that the work on the DTN-related protocols will have growing importance to terrestrial communications, and we are seeking partnerships with developers of such systems in order to widen our pool of expertise.

For interplanetary applications we plan to demonstrate the potential evolution of the current "CCSDS File Delivery Protocol" (CFDP), which is a delay tolerant file transfer protocol, towards utilization of Bundling.

CFDP

The CFDP enables the moving of a file from one filestore to another, where the two filestores are in general resident in separate data systems and usually with an intervening space link. In addition to the purely file delivery-related functions, the protocol also includes file management services to allow control over the storage medium. It is a delay tolerant protocol whose model of operations is much like e-mail that conveys files as attachments. The protocol as currently designed contains its own reliability mechanisms and does not assume an underlying retransmission capability. The CFDP concept includes its own notion of "custodial transfer" where a sender can transmit a file to a receiver over a single link and, upon receipt of the entire file, the receiver can notify the sender that it will take care of any successive forward transmission hops. This allows the sender to release local processing and storage resources and to deploy them on new data acquisition – a very important feature for transmission of data to or from nodes with limited resources in networks with long signal propagation delays.

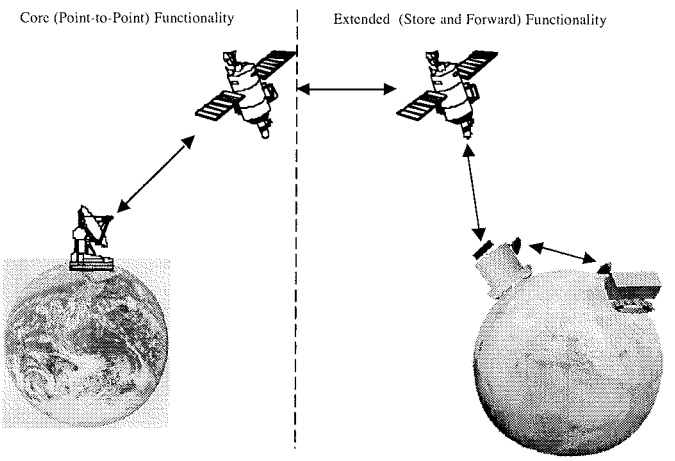

In its simplest form, the protocol provides a *Core* file delivery capability operating across a single link. For more complex mission scenarios, the protocol offers *Extended* operation providing store-and-forward functionality across an arbitrary network, containing multiple links with disparate availability, as well as subnetworks with heterogeneous protocols. The protocol can operate over a wide range of underlying communication services. The Core procedures constitute the interaction between two protocol entities with a direct network path between them. Where direct network connectivity between the source and destination is impossible, the Extended procedures automatically build an end-to-end file copy transaction by executing multiple file copy operations, as follows: one file copy operation between the source and the first waypoint; others between successive waypoints as necessary; and a final file copy operation between the last waypoint and the destination. Each of these is simply another instance of the Core file copy operation. The reliability of a transaction is determined by whether the transaction is chosen to operate in unacknowledged mode or in one of the acknowledged modes. In unacknowledged mode data delivery failures are not reported to the sender and, therefore, cannot be repaired, although errors will be detected and erroneous data discarded. Reception of the complete file is therefore not guaranteed. In acknowledged mode, the receiver informs the sender of any undelivered file segments or ancillary data. These are then retransmitted, guaranteeing complete file delivery.

The CFDP is tailored to the needs of space faring vehicles. It is fairly complex, and it is not of great interest outside of the space oriented community. It does have the potential of operating well over the Bundling protocol (that is, utilizing Bundling as its underlying communication service), thus extending its useful life and perhaps broadening its community of interest. The CFDP is intentionally to some extent a prototype of some major features of the DTN architecture. The experience gained with it in the development particularly of custody transfer procedures has been very valuable.

Bundling Protocol

The Bundling protocol suite is the vehicle by which the stitching together of the time-disjoint individual hops in a DTN occurs. Interconnection of adjacent nodes may occur nearly instantaneously, or the adjacent nodes may be separated by variable and arbitrarily long time delays. Nodes at the termini of the hops operate in a store and forward mode, and many of them may also be "custodial" - they will accept responsibility for

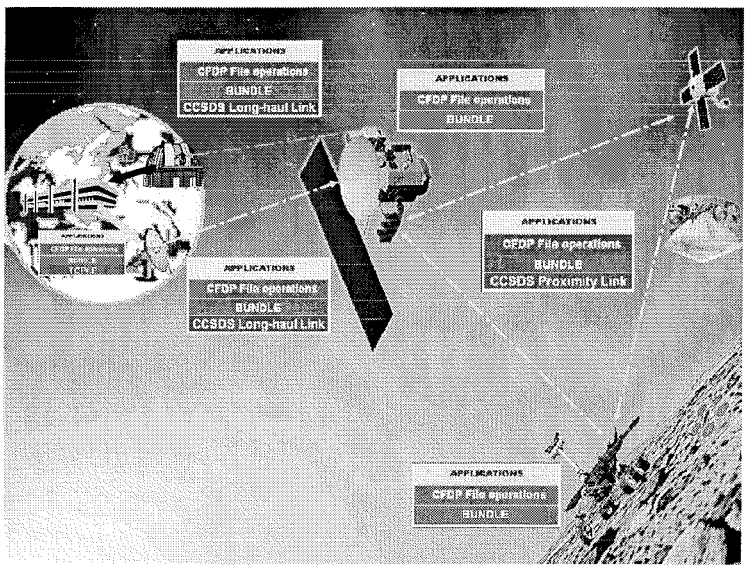

relaying an incoming Bundle, thus releasing a prior node from this responsibility and liberating its processing, storage and communications resources. By using such techniques, a highly robust network may be built that is tolerant of both delay and disconnectedness. Note that delay can be introduced by several factors, including a lack of bi-directional communications capability that makes it impossible to verify that data have been reliably delivered in real time.

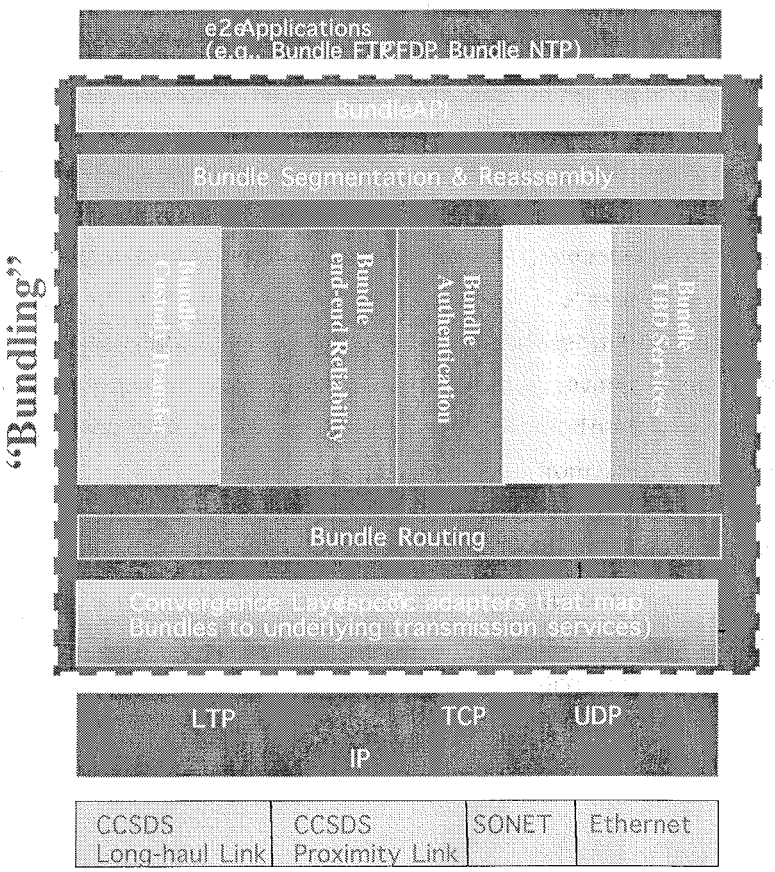

The DTN architecture differs from CFDP in several key respects:

DTN is not confined to supporting just file transfer; it can handle virtually any end-to-end application across any concatenation of different environments. Eventually, we think, CFDP's file delivery and filestore management services operating in Core unacknowledged mode will simply be one of those applications, with responsibility for multi-hop forwarding and reliable transmission delegated to Bundling.

DTN's internal functions are more clearly layered than CFDP, so it should be easier to evolve.

Bundling will provide a rich and comprehensive set of network services, including flexible and dynamic routing, comprehensive security measures, and congestion control.

Evolution and Strategy

We have begun the mainstream development of the Bundling protocol. We are committed to use experimentation and prototyping as an integral part of this development lifecycle. We are using this prototyping experience to develop draft written specifications of the key protocols, which along with working code may be shared with others so that they may themselves implement prototype Bundle nodes. For interplanetary applications we plan to demonstrate the potential evolution of the current "CCSDS File Delivery Protocol" towards utilization of Bundling.

Tactically, it is important to foster widespread use of CFDP as a prototype for DTN concepts, and to avoid increasing the complexity of CFDP beyond its present state, allowing Bundling to handle the

complexity required by DTNs. This should then be followed by repositioning CFDP as an application running over Bundling.

It is highly desirable to mobilize DTN users to develop Bundling as a community effort. This would result in faster development and a more robust protocol. In addition, broadening of the base of users and applications would enhance and speed up the deployment of Bundling into the field.

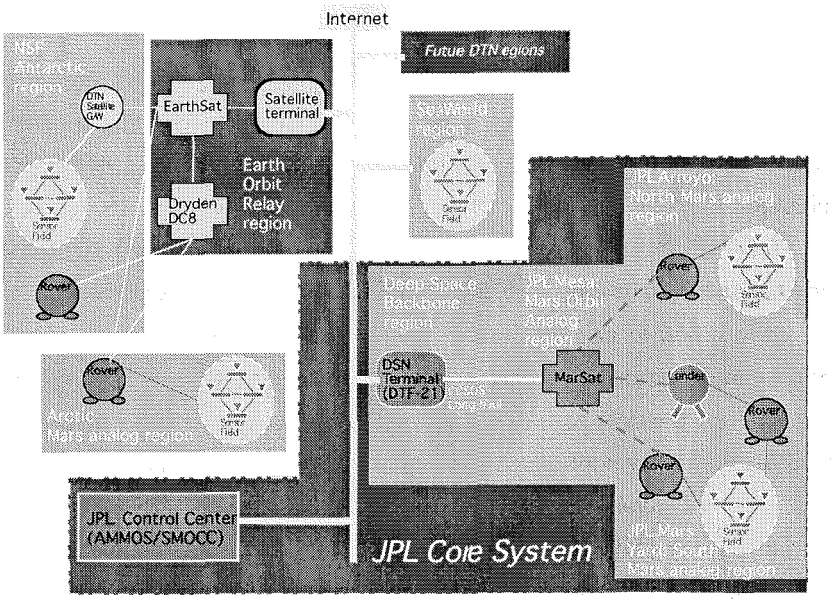

Strategically, to achieve the goal of implementing and deploying Bundling specifically, and implementing the Interplanetary Internet generally, we must obtain a shared commitment among the space, commercial, and military communities to make it happen. We need to create a test and demonstration environment that shows the progressive maturity of DTN technology. And perhaps most importantly, it is essential that we gain a commitment of organizational resources that recognizes the pivotal role of this work for the rest of this decade.

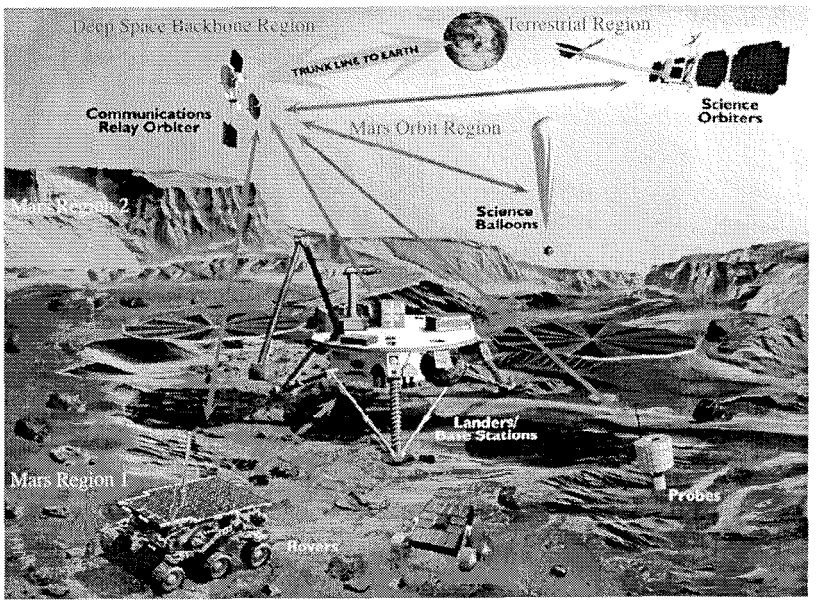

FINAL

CL 04-914

MP

Proposed Evolution of CCSDS Protocols - Strategy and Technology

Richard D. Carper[1], Scott Burleigh[2], Adrian J. Hooke[2], J. Leigh Torgerson[2]

[1]*Consultant, Space Data Systems, Corvallis, Oregon.*

[2]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California.

Background

For more than two decades the international Consultative Committee for Space Data Systems (CCSDS) has developed and promulgated space link standards for the use of the international space-faring community. Nearly every member of that community is a member of the CCSDS. The organization has been sufficiently effective in its work that at the present time over 250 spacecraft have used or are using these standards. Using platforms range from high altitude balloons through all sizes and types of earth orbiters, including the International Space Station, through interplanetary spacecraft and planetary landers.

Currently some of the CCSDS member Agencies, most specifically the Jet Propulsion Laboratory (JPL) of the California Institute of Technology, are involved in the development of standards which allow the extension of the Internet into the space data domain. That domain consists of multiple regions which constitute a disconnected, variable-delay environment, and there is a growing need for a general, standard way to communicate end-to-end through this environment. The fundamental concept which JPL is using is that local in-situ short delay Internets, distributed across the Solar System on free flying spacecraft and on and around other planets, are interconnected via a long delay deep space backbone network. Just as the TCP/IP suite unites the Earth's Internet as a "network of networks", a new protocol suite called "Bundling" unites the InterPlaNetary Internet (IPN) into a "network of Internets".

JPL is testing the architectural concepts by beginning the detailed definition and prototyping of Bundling, and has charted a strategy whereby today's space systems can evolve smoothly into the Bundling era throughout the coming decade. It is especially noteworthy that the scope of the IPN has also been broadened so that it is a member of a new class of "Disruption Tolerant Networks" (DTNs).[1]

Disruption Tolerant Networks

In many environments – such as interplanetary communications – a real time end-to-end path between two users may never exist as a connected entity. Instead, communications must occur by

[1] This work was performed in part at and under the direction of the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration.

stitching together a series of time-disjoint individual hops. The time-disjoint character of point-to-point connections in a DTN may range from known, predicted losses of line-of-sight contact between orbiting vehicles and planetary rovers, through the unpredictable change of location of individual mobile units, to pre-emption of resources in a battlefield situation or destruction of one or more intermediate nodes in the network. While deep space communications (with their enormous signal propagation delays) are obvious examples of DTNs, we expect that this mode of operation will become increasingly important to terrestrial communications as well. It seems evident that with the tremendous increase in the use of pocket size communicating devices of many different kinds and of similar vehicular devices, the networks which support them will inherently be DTNs. The result of this Internet evolution is that user devices are becoming increasingly mobile, and the end-to-end data path is traveling across a wider variety of environments. Mobility coupled with power, weight, volume constraints suggests that end-to-end communications dialogs across time disjoint connectivity are going to be highly likely, if not inevitable. The DTN protocol mechanisms could significantly enhance the Internet's ability to accommodate the extreme heterogeneity that is emerging, by bridging between islands ("regions") of homogeneity to support end-to-end delivery of information.

Therefore we believe that the work on the DTN-related protocols will have growing importance to terrestrial communications, and we are seeking partnerships with developers of such systems in order to widen our pool of expertise.

For interplanetary applications we plan to demonstrate the potential evolution of the current "CCSDS File Delivery Protocol" (CFDP), which is a delay tolerant file transfer protocol, towards utilization of Bundling.

CFDP

The CFDP enables the moving of a file from one filestore to another, where the two filestores are in general resident in separate data systems and usually with an intervening space link. In addition to the purely file delivery-related functions, the protocol also includes file management services to allow control over the storage medium. It is a delay tolerant protocol whose model of operations is much like e-mail that conveys files as attachments. The protocol as currently designed contains its own reliability mechanisms and does not assume an underlying retransmission capability. The CFDP concept includes its own notion of "custodial transfer" where a sender can transmit a file to a receiver over a single link and, upon receipt of the entire file, the receiver can notify the sender that it will take care of any

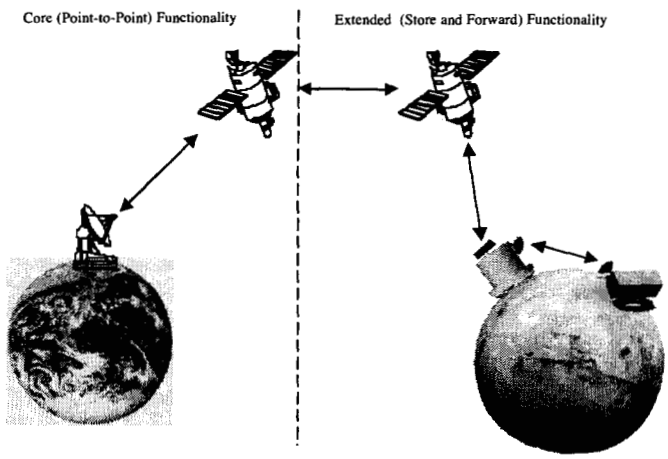

successive forward transmission hops. This allows the sender to release local processing and storage resources and to deploy them on new data acquisition – a very important feature for transmission of data to or from nodes with limited resources in networks with long signal propagation delays.

In its simplest form, the protocol provides a *Core* file delivery capability operating across a single link. For more complex mission scenarios, the protocol offers *Extended* operation providing store-and-forward functionality across an arbitrary network, containing multiple links with disparate availability, as well as subnetworks with heterogeneous protocols. The protocol can operate over a wide range of underlying communication services. The Core procedures constitute the interaction between two protocol entities with a direct network path between them. Where direct network connectivity between the source and destination is impossible, the Extended procedures automatically build an end-to-end file copy transaction by executing multiple file copy operations, as follows: one file copy operation between the source and the first waypoint; others between successive waypoints as necessary; and a final file copy operation between the last waypoint and the destination. Each of these is simply another instance of the Core file copy operation. The reliability of a transaction is determined by whether the transaction is chosen to operate in unacknowledged mode or in one of the acknowledged modes. In unacknowledged mode data delivery failures are not reported to the sender and, therefore, cannot be repaired, although errors will be detected and erroneous data discarded. Reception of the complete file is therefore not guaranteed. In acknowledged mode, the receiver informs the sender of any undelivered file segments or ancillary data. These are then retransmitted, guaranteeing complete file delivery.

The CFDP is tailored to the needs of space faring vehicles. It is fairly complex, and it is not of great interest outside of the space oriented community. It does have the potential of operating well over the Bundling protocol (that is, utilizing Bundling as its underlying communication service), thus extending its useful life and perhaps broadening its community of interest. The CFDP is intentionally to some extent a prototype of some major features of the DTN architecture. The experience gained with it in the development particularly of custody transfer procedures has been very valuable.

Bundling Protocol

The Bundling protocol suite is the vehicle by which the stitching together of the time-disjoint individual hops in a DTN occurs. Interconnection of adjacent nodes may occur nearly instantaneously, or the adjacent nodes may be separated by variable and arbitrarily long time delays. Nodes at the termini of the hops operate in a store and

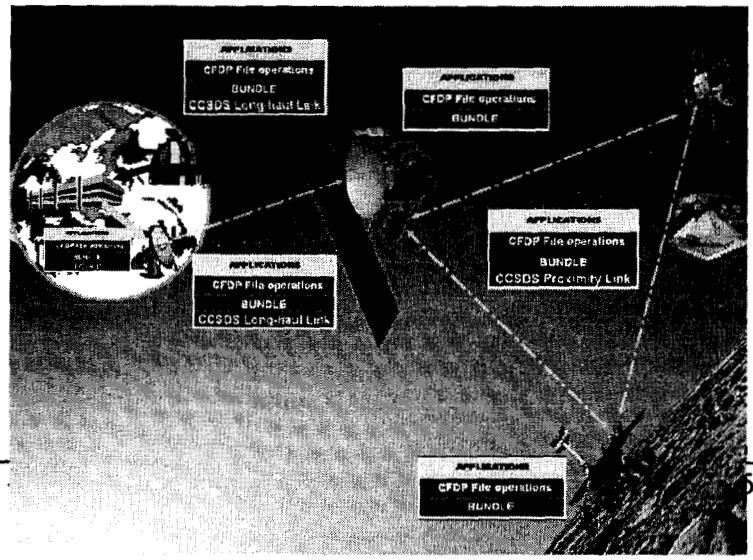

forward mode, and many of them may also be "custodial" - they will accept responsibility for relaying an incoming Bundle, thus releasing a prior node from this responsibility and liberating its processing, storage and communications resources. By using such techniques, a highly robust network may be built that is tolerant of both delay and disconnectedness. Note that delay can be introduced by several factors, including a lack of bi-directional communications capability that makes it impossible to verify that data have been reliably delivered in real time.

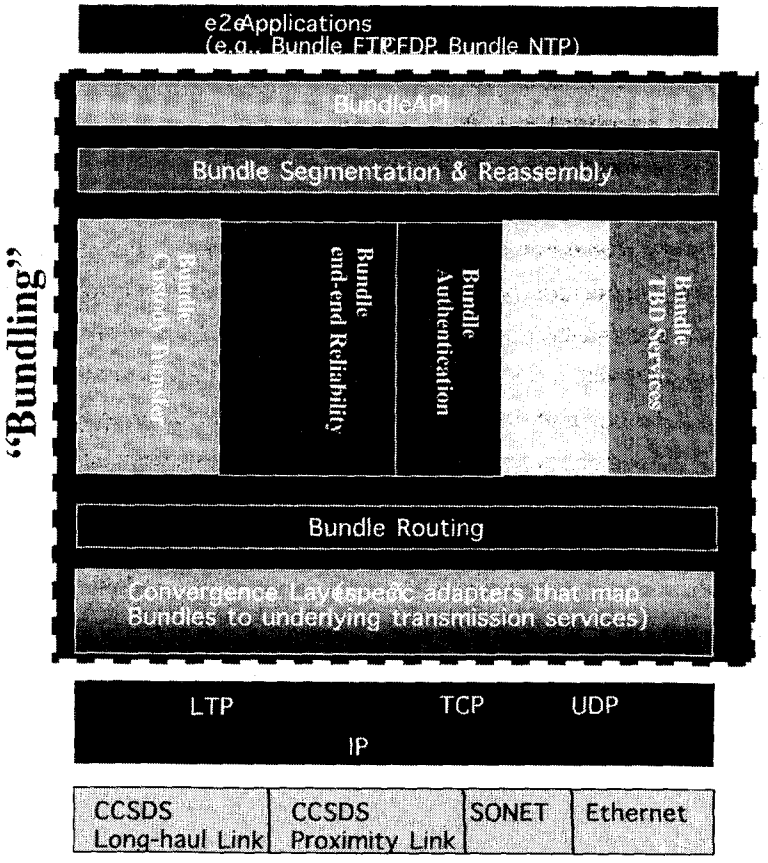

The DTN architecture differs from CFDP in several key respects:

DTN is not confined to supporting just file transfer; it can handle virtually any end-to-end application across any concatenation of different environments. Eventually, we think, CFDP's file delivery and filestore management services operating in Core unacknowledged mode will simply be one of those applications, with responsibility for multi-hop forwarding and reliable transmission delegated to Bundling.

DTN's internal functions are more clearly layered than CFDP, so it should be easier to evolve.

Bundling will provide a rich and comprehensive set of network services, including flexible and dynamic routing, comprehensive security measures, and congestion control.

Evolution and Strategy

We have begun the mainstream development of the Bundling protocol. We are committed to use experimentation and prototyping as an integral part of this development lifecycle. We are using this prototyping experience to develop draft written specifications of the key protocols, which along with working code may be shared with others so that they may themselves implement prototype Bundle nodes. For interplanetary applications we plan to demonstrate the potential evolution of the current "CCSDS File Delivery Protocol" towards utilization of Bundling.

Tactically, it is important to foster widespread use of CFDP as a prototype for DTN concepts, and

to avoid increasing the complexity of CFDP beyond its present state, allowing Bundling to handle the complexity required by DTNs. This should then be followed by repositioning CFDP as an application running over Bundling.

It is highly desirable to mobilize DTN users to develop Bundling as a community effort. This would result in faster development and a more robust protocol. In addition, broadening of the base of users and applications would enhance and speed up the deployment of Bundling into the field.

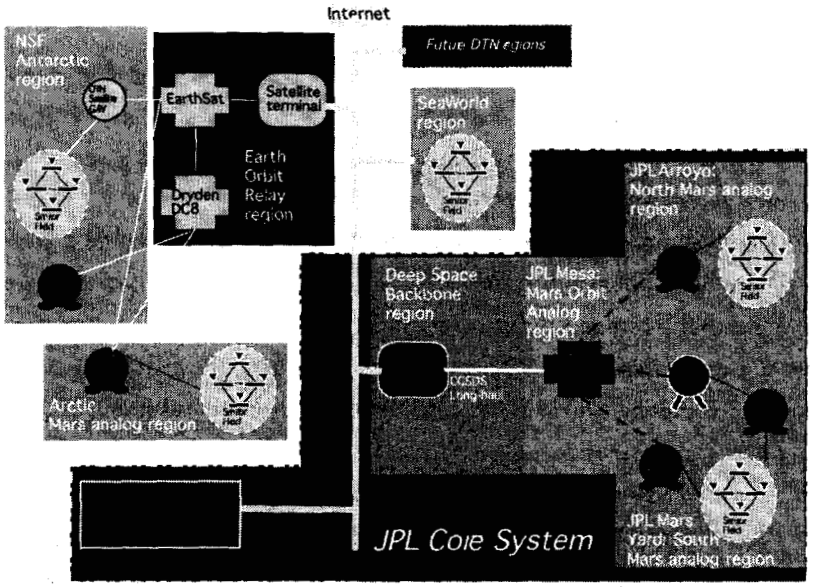

Strategically, to achieve the goal of implementing and deploying Bundling specifically, and implementing the Interplanetary Internet generally, we must obtain a shared commitment among the space, commercial, and military communities to make it happen. We need to create a test and demonstration environment that shows the progressive maturity of DTN technology. And perhaps most importantly, it is essential that we gain a commitment of organizational resources that recognizes the pivotal role of this work for the rest of this decade.