The Consultative Committee for Space Data Systems

Report Concerning Space Data System Standards

PROXIMITY-1 SPACE LINK PROTOCOL—

RATIONALE, ARCHITECTURE, AND SCENARIOS

INFORMATIONAL REPORT

CCSDS 210.0-G-1

GREEN BOOK
August 2007

The Consultative Committee for Space Data Systems

Report Concerning Space Data System Standards

PROXIMITY-1 SPACE LINK PROTOCOL—

RATIONALE, ARCHITECTURE, AND SCENARIOS

INFORMATIONAL REPORT

CCSDS 210.0-G-1

GREEN BOOK

August 2007

AUTHORITY

Issue:	Green Book, Issue 1
Date:	August 2007
Location:	Washington, DC, USA

This document has been approved for publication by the Management Council of the Consultative Committee for Space Data Systems (CCSDS) and reflects the consensus of technical working group experts from CCSDS Member Agencies. The procedure for review and authorization of CCSDS Reports is detailed in the *Procedures Manual for the Consultative Committee for Space Data Systems*.

This document is published and maintained by:

> CCSDS Secretariat
> Office of Space Communication (Code M-3)
> National Aeronautics and Space Administration
> Washington, DC 20546, USA

FOREWORD

This document is a CCSDS Report which contains background and explanatory material to support the CCSDS Recommended Standards for Proximity-1 Space Link Protocol (references [8]-[10]).

Through the process of normal evolution, it is expected that expansion, deletion, or modification of this document may occur. This Report is therefore subject to CCSDS document management and change control procedures, which are defined in the *Procedures Manual for the Consultative Committee for Space Data Systems*. Current versions of CCSDS documents are maintained at the CCSDS Web site:

http://www.ccsds.org/

Questions relating to the contents or status of this document should be addressed to the CCSDS Secretariat at the address indicated on page i.

At time of publication, the active Member and Observer Agencies of the CCSDS were:

Member Agencies

- – Agenzia Spaziale Italiana (ASI)/Italy.
- – British National Space Centre (BNSC)/United Kingdom.
- – Canadian Space Agency (CSA)/Canada.
- – Centre National d'Etudes Spatiales (CNES)/France.
- – Deutsches Zentrum für Luft- und Raumfahrt e.V. (DLR)/Germany.
- – European Space Agency (ESA)/Europe.
- – Federal Space Agency (FSA)/Russian Federation.
- – Instituto Nacional de Pesquisas Espaciais (INPE)/Brazil.
- – Japan Aerospace Exploration Agency (JAXA)/Japan.
- – National Aeronautics and Space Administration (NASA)/USA.

Observer Agencies

- – Austrian Space Agency (ASA)/Austria.
- – Belgian Federal Science Policy Office (BFSPO)/Belgium.
- – Central Research Institute of Machine Building (TsNIIMash)/Russian Federation.
- – Centro Tecnico Aeroespacial (CTA)/Brazil.
- – Chinese Academy of Sciences (CAS)/China.
- – Chinese Academy of Space Technology (CAST)/China.
- – Commonwealth Scientific and Industrial Research Organization (CSIRO)/Australia.
- – Danish National Space Center (DNSC)/Denmark.
- – European Organization for the Exploitation of Meteorological Satellites (EUMETSAT)/Europe.
- – European Telecommunications Satellite Organization (EUTELSAT)/Europe.
- – Hellenic National Space Committee (HNSC)/Greece.
- – Indian Space Research Organization (ISRO)/India.
- – Institute of Space Research (IKI)/Russian Federation.
- – KFKI Research Institute for Particle & Nuclear Physics (KFKI)/Hungary.
- – Korea Aerospace Research Institute (KARI)/Korea.
- – MIKOMTEK: CSIR (CSIR)/Republic of South Africa.
- – Ministry of Communications (MOC)/Israel.
- – National Institute of Information and Communications Technology (NICT)/Japan.
- – National Oceanic and Atmospheric Administration (NOAA)/USA.
- – National Space Organization (NSPO)/Taiwan.
- – Naval Center for Space Technology (NCST)/USA.
- – Space and Upper Atmosphere Research Commission (SUPARCO)/Pakistan.
- – Swedish Space Corporation (SSC)/Sweden.
- – United States Geological Survey (USGS)/USA.

DOCUMENT CONTROL

Document	Title	Date	Status
CCSDS 210.0-G-1	Proximity-1 Space Link Protocol—Rationale, Architecture, and Scenarios, Informational Report, Issue 1	August 2007	Current issue

CONTENTS

CONTENTS (continued)

1 INTRODUCTION

1.1 PURPOSE AND SCOPE

This document is an adjunct document to the three CCSDS Recommended Standards for Proximity-1 Space Link Protocol: Data Link Layer (reference [8]), Synchronization and Channel Coding (reference [9]), and Physical Layer (reference [10]). It contains material helpful in understanding these Recommended Standards, which will assist decision makers and implementers in evaluating the applicability of the protocol to mission needs and in making implementation, option selection, and configuration decisions related to the protocols.

This report provides supporting descriptive and tutorial material. **This document is not part of the Recommended Standards.** In the event of conflicts between this report and the Recommended Standards, the Recommended Standards shall prevail.

1.2 STRUCTURE OF THIS DOCUMENT

This document is organized as follows.

- Section 1 defines the purpose and scope of this document, and lists the definition and convention used throughout this document.

- Section 2 gives an overview of the application context and concept and rationale of the Proximity-1 protocol.

- Section 3 describes the protocol architecture and the functions and services provided by each protocol entity.

- Section 4 describes the data unit structures and discusses various aspects of the operation of the protocol.

- Section 5 provides several possible scenarios under which the Proximity-1 Link Protocol may be operated.

- Section 6 provides background information about the performance trades made during the development of the protocol.

1.3 DEFINITIONS

1.3.1 GENERAL

This Report makes use of a number of terms defined in the OSI Reference Model and the OSI Service Model (references [1] and [2]) to describe, in a generic sense, the technology and the services provided for information exchange between two systems. These terms include, but are not exclusive to, blocking, connection, Data Link layer, entity, flow control,

Network layer, peer entities, Physical layer, protocol control information, Protocol Data Unit (PDU), real system, segmenting, service, Service Access Point (SAP), SAP address, Service Data Unit (SDU), confirmation, indication, primitive, request, response, service provider, and service user.

1.3.2 TERMS

The following terms are defined in this Report. Note that new terms are added and some of the definitions are modified for clarity.

asynchronous channel: a data channel where the symbol data are modulated onto the channel only for the period of the message. The message must be preceded by an acquisition sequence to achieve symbol synchronization. Bit synchronization must be reacquired on every message. A hailing channel is an example of an asynchronous channel.

asynchronous data link: a data link consisting of a sequence of variable-length Proximity Link Transmission Units (PLTUs), which are not necessarily concatenated. Two types of asynchronous data links are:

1) Asynchronous Data Link over an Asynchronous Channel

   Hailing provides an example of an asynchronous data link over an asynchronous channel. An important issue is resynchronization between successive hails. Idle is provided for the reacquisition process.

2) Asynchronous Data Link over a Synchronous Channel

   Data service provides an example of an asynchronous data link over a synchronous channel. Once the link is established via hailing, communication transitions to a synchronous channel and maintains the link in this configuration until the session is interrupted or ends. If the Physical layer does not receive data from the Data Link layer, it provides idle to maintain a synchronous channel.

caller and responder: A **caller transceiver** is the initiator of the link establishment process and manager of negotiation (if required) of the session. A **responder transceiver** typically receives link establishment parameters from the caller. The caller initiates communication between itself and a responder on a pre-arranged communications channel with predefined controlling parameters. As necessary, the caller and responder may negotiate the controlling parameters for the session (at some level between fully controlled and completely adaptive).

COP-P: Communication Operations Procedure for Proximity links (COP-P). The COP-P includes both the FARM-P and FOP-P of the caller and responder unit.

FARM-P: Frame Acceptance and Reporting Mechanism for Proximity links, for Sequence Controlled service carried out within the receiver in the Proximity-1 link.

FOP-P: Frame Operation Procedure for Proximity links for ordering the output frames for Sequence Controlled service carried out in the transmitter in the Proximity-1 link.

forward link: that portion of a Proximity space link in which the caller transmits and the responder receives (typically a command link).

hailing: the persistent activity used to establish a Proximity link by a caller to a responder in either full or half duplex. It does not apply to simplex operations.

hailing channel: the forward and return frequency pairs that a caller and responder use to establish physical link communications.

mission phase: a mission period during which specified communications characteristics are fixed. The transition between two consecutive mission phases may cause an interruption of the communications services.

PCID: The Physical Channel ID is used to distinguish between Proximity Link Control Words (PLCWs) received on a single receive channel in support of two independent transmitting channels.

P-frame: a Version-3 Transfer Frame that contains only self-identified and self-delimited supervisory protocol data units; compare U-frame.

physical channel: The RF channel upon which the stream of bits is transferred over a space link in a single direction.

PLCW: Proximity Link Control Word. The PLCW is the protocol data unit for reporting Sequence Controlled service status via the return link from the responder back to the caller.

PLTU: Proximity Link Transmission Unit. The PLTU is the data unit composed of the Attached Synchronization Marker, the Version-3 Transfer Frame, and the attached Cyclic Redundancy Check (CRC)-32.

Protocol object: directives, PLCWs, or status reports contained within an SPDU.

Proximity link: short-range, bi-directional, fixed or mobile radio links, generally used to communicate among probes, landers, rovers, orbiting constellations, and orbiting relays. These links are characterized by short time delays, moderate (not weak) signals, and short, independent sessions.

pseudo packet ID: the temporary packet ID assigned by the protocol to a user's packet within the segmentation process.

resynchronization (COP-P): process in which sender and receiver nodes readjust their Sequence Controlled frame numbers via the SET V(R) activity.

return link: that portion of a Proximity space link in which the responder transmits and the caller receives (typically a telemetry link).

Routing ID: identifier that uniquely identifies a user's packet through the segmentation process. It consists of a PCID, Port ID, and pseudo packet ID.

Sent queue (Sent Frame queue): A container for Sequence Controlled frames that have been sent but not yet acknowledged by the receiver.

session: a continuous dialog between two communicating Proximity link transceivers. It consists of three distinct operational phases: session establishment, data services, and session termination.

space link: a communications link between transmitting and receiving entities, at least one of which is in space.

SPDU: Supervisory Protocol Data Unit. The SPDU is used by the local transceiver either to control or to report status to the remote partnered transceiver. It consists of one or more directives, reports, or PLCWs.

synchronous channel: a data channel where the symbol data are continuously modulated onto the channel at a fixed data rate. If the data link fails to provide frames (data or fill), it is the responsibility of the Physical layer to provide the continuous bit stream.

U-frame: a Version-3 Transfer Frame that contains user data information; compare P-frame.

vehicle controller: the entity (e.g., spacecraft control computer) which receives the notifications defined in annex D of reference [8] and potentially acts upon them.

Version-3 Transfer Frame: a Proximity-1 transfer frame.

1.4   REFERENCES

The following documents are referenced in this Report. At the time of publication, the editions indicated were valid. All documents are subject to revision, and users of this Report are encouraged to investigate the possibility of applying the most recent editions of the documents indicated below. The CCSDS Secretariat maintains a register of currently valid CCSDS documents.

[1] *Information Technology—Open Systems Interconnection—Basic Reference Model: The Basic Model*. International Standard, ISO/IEC 7498-1:1994. 2nd ed. Geneva: ISO, 1994.

[2] *Information Technology—Open Systems Interconnection—Basic Reference Model— Conventions for the Definition of OSI Services*. International Standard, ISO/IEC 10731:1994. Geneva: ISO, 1994.

[3]   *Communications Operation Procedure-1*. Recommendation for Space Data System Standards, CCSDS 232.1-B-1. Blue Book. Issue 1. Washington, D.C.: CCSDS, September 2003.

[4]   *TC Space Data Link Protocol*. Recommendation for Space Data System Standards, CCSDS 232.0-B-1. Blue Book. Issue 1. Washington, D.C.: CCSDS, September 2003.

[5]   *Space Packet Protocol*. Recommendation for Space Data System Standards, CCSDS 133.0-B-1. Blue Book. Issue 1. Washington, D.C.: CCSDS, September 2003.

[6]   *TC Synchronization and Channel Coding*. Recommendation for Space Data System Standards, CCSDS 231.0-B-1. Blue Book. Issue 1. Washington, D.C.: CCSDS, September 2003.

[7]   *CCSDS Global Spacecraft Identification Field Code Assignment Control Procedures*. Recommendation for Space Data System Standards, CCSDS 320.0-B-4. Blue Book. Issue 4. Washington, D.C.: CCSDS, January 2006.

[8]   *Proximity-1 Space Link Protocol—Data Link Layer*. Recommendation for Space Data System Standards, CCSDS 211.0-B-4. Blue Book. Issue 4. Washington, D.C.: CCSDS, July 2006.

[9]   *Proximity-1 Space Link Protocol—Coding and Synchronization Sublayer*. Recommendation for Space Data System Standards, CCSDS 211.2-B-1. Blue Book. Issue 1. Washington, D.C.: CCSDS, April 2003.

[10]   *Proximity-1 Space Link Protocol—Physical Layer*. Recommendation for Space Data System Standards, CCSDS 211.1-B-3. Blue Book. Issue 3. Washington, D.C.: CCSDS, March 2006.

[11]   *AOS Space Data Link Protocol*. Recommendation for Space Data System Standards, CCSDS 732.0-B-2. Blue Book. Issue 2. Washington, D.C.: CCSDS, July 2006.

[12]   *TM Space Data Link Protocol*. Recommendation for Space Data System Standards, CCSDS 132.0-B-1. Blue Book. Issue 1. Washington, D.C.: CCSDS, September 2003.

[13]   *Encapsulation Service*. Recommendation for Space Data System Standards, CCSDS 133.1-B-1. Blue Book. Issue 1. Washington, D.C.: CCSDS, June 2006.

[14]   *Space Link Identifiers*. Recommendation for Space Data System Standards, CCSDS 135.0-B-3. Blue Book. Issue 3. Washington, D.C.: CCSDS, October 2006.

2 OVERVIEW

2.1 INTRODUCTION

The Proximity-1 protocol is a 'Space Link Protocol' designed for the purpose of proximate communications among 'probes, landers, rovers, orbiting constellations, and orbiting relays' (reference [8]).

To aid the efficient design of the protocol, an application domain and a prototypical application within that domain were identified (as described below). This approach has provided the protocol design team with a focal point for discussion, and it has provided feedback that has assisted with the refinement and further development of the protocol. The protocol is still undergoing development so it will be able to service the needs and requirements of future missions in a variety of operational scenarios.

2.2 APPLICATION DOMAIN OF THE PROXIMITY-1 PROTOCOL

2.2.1 GENERAL

The Application Domain for the Proximity-1 protocol is communications among 'probes, landers, rovers, orbiting constellations, and orbiting relays' (reference [8]).

The Application Profile for the development of Proximity-1 has been wireless communication between space assets whose inter-spacecraft distances range between ~1 meter and approximately 100,000 kilometers, though greater distances could be accommodated. Example applications include (but are not limited to):

– a micro-rover-to-lander scenario (involving ranges as short as one meter);

– landed assets on another planet with a relay orbiter;

– landed assets with a spacecraft approaching the planet.

The name 'Proximity' contrasts this protocol to the CCSDS Packet Telemetry/ AOS/Telecommand protocols (see references [12], [11], and [4], respectively) used for Earth-deep space links having extremely long communication distances.

Communication distance, however, was not intended to be the only distinguishing feature. With Promixity-1, there is no Ground operator manual intervention, and resources such as processing power, storage, etc., are usually limited at both ends of the link.

Currently data rates of 2.048 kilobits per second are defined; thus modest- to low-delay bandwidth products are involved.

2.2.2   THE PROTOTYPICAL APPLICATION

To provide a focus for the development of the Proximity-1 Protocol and to help identify the requirements driving its definition, a generic prototype application was identified. This application has largely been realized in the 2004-2005 Mars missions that have demonstrated the effectiveness of the protocol.

While this application does not explore the full potential of Proximity-1 applications, it does provide an illustration of the core capabilities. It is used in this document to aid description. Further discussion of the full range of features appears in later sections of this document.

The Prototypical Application involves a Relay Orbiter and Surface Assets (e.g., rovers) operating on a planet/satellite. The Proximity-1 protocol has the capability of achieving reliable communication (i.e., data transferred in order without duplicates or gaps) within a communication session between the Relay Orbiter and the Rovers. These communications must be able to be conducted without real-time intervention from Earth and provide appropriate levels of data throughput and link performance.

The role of each element in this scenario can be expressed in three parts:

1) Orbiter establishes communication with the Rovers by hailing and relays/forwards commands sent from Earth and typically received by the Orbiter before the over flight. The forward Proximity-1 link can be one-to-many (see figure 2-1a).
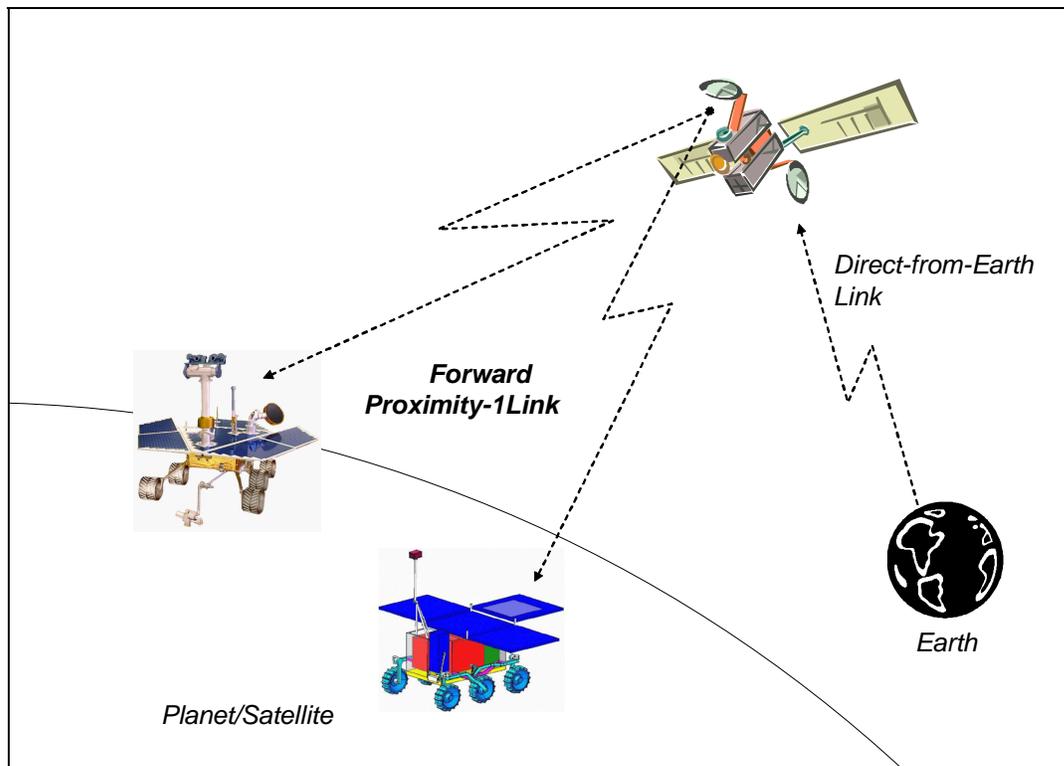
Figure 2-1a: Relay Command Link

2) Each Rover responds separately to the hail by telemetering data to the Orbiter to be relayed/returned to Earth when a direct-to-Earth link becomes available. Note forwarding commands and receiving telemetry over the Proximity link can be done concurrently (see figure 2-1b).
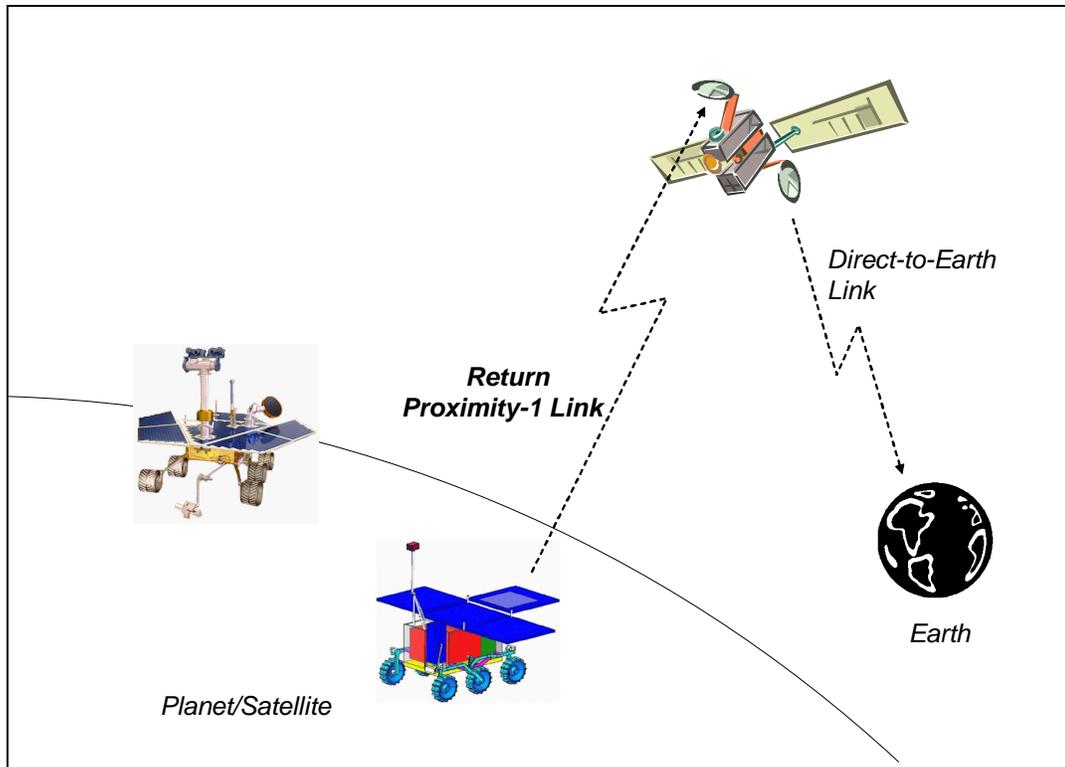
Figure 2-1b: Relay Telemetry Link

Orbiter provides additional services (time correlation, time transfer, resource sharing) enabled via Proximity-1 messaging between orbiter and rover applications. Note that relaying is not a required function of Proximity-1, and in particular the data exchanged across the Proximity-1 data link could be consumed (in the 'final destination' sense) by application entities at both endpoints.
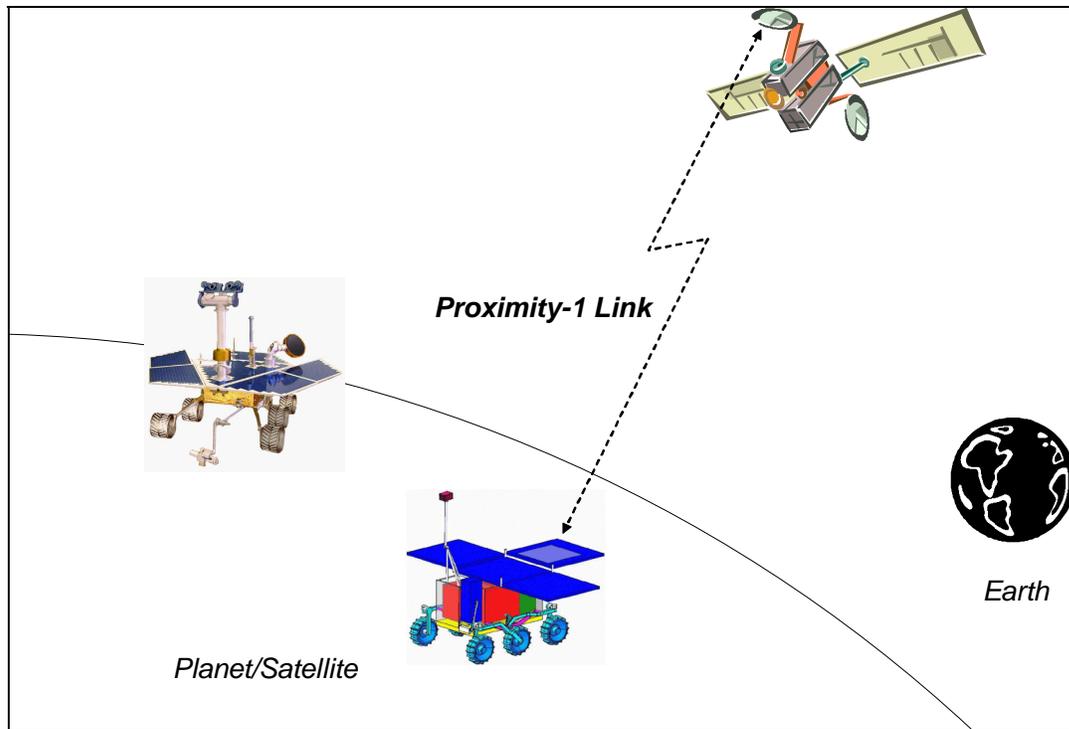
Figure 2-1c: Proximity-1 Messaging

The Proximity-1 Recommended Standards facilitate a standard approach to messaging by assigning well-known Port IDs (reference [14]) to Proximity message types. On the forward link (Orbit to Surface), hardware commands are assigned to Port ID 1, so that once validated they can be routed directly from the Proximity transceiver to the addressed hardware unit minimizing latency and bypassing the flight computer. For both forward and return links (surface to orbit), CCSDS Space Packets are assigned to Port ID 2, so that these self-delimiting data units, i.e., packets containing application data (data, status, or control messages), can be transferred and immediately interpreted by the receiving node. Well-defined Port IDs for both the forward and return link are assigned to specific data types in the Space Link Identifiers Recommended Standard (reference [14]).

Although simplex operations are part of the protocol, the primary benefits of Proximity-1 arise with two-way operations, full duplex or half duplex. In continuing with the orbiter-rover scenario to aid description of Proximity-1, a two-way operation using either full duplex or half duplex is assumed.

Several key drivers arise from this prototypical application as a result of its being used in a space mission and include:

− predictable episodic connectivity;

− access scheme;

− traffic characteristics of space science communications;

–   asymmetry of communications resources;

–   large variety of user missions;

–   efficient use of resources;

–   appropriate balance of manual vs. autonomous operations.

While this relay orbiter application is identified as particularly suitable for Proxmity-1, it is emphasized that the onward relaying of data (e.g., to/from Earth and to/from the orbiter) is *not* part of Proximity-1 itself (this represents a higher-layer function).

Each of the listed drivers is described in the following paragraphs:

a) **Predictably Episodic Connectivity** – There is a dynamic but nonetheless highly predictable connectivity; i.e., times when the physical links will be available are deterministic. Note that when no link is available, transceiver resources may be conserved (see item 5 below). Predictability of connection times allows ground operations to schedule each contact through either sequencing the Orbiter and asset separately or using the capability of Proximity-1 to enable autonomous link establishment through hailing. Either the Orbiter or the asset could be programmed to establish the link based upon over-flight geometry and conduct data operations until both sides negotiate the termination of the link via Proximity-1 link termination procedures.

b) **Access Scheme** – There will not be a large number of surface assets (landers/rovers) in view at once (at least for the foreseeable future) and the need is for point-to-point communication on the return link and at most one-to-many on the forward link. If there are multiple assets in view of the Orbiter at once, the capability of providing low latency access among them applies differently than in conventional multi-access schemes used on Earth. The reason for this is celestial mechanics. Depending upon the geometry, the orbiter typically cannot view the Earth and these assets simultaneously. Therefore, the orbiter must first store the data it receives during an over flight and later forward that data when the earth is in view. As long as adequate bandwidth is provided to each asset during the over flight, then each asset can be adequately serviced by the Orbiter. Thus a series of time-sequenced point-to-point connections is preferable to multi-access protocols, because of the greater efficiencies in throughput and simplicity of time sequencing the link.

c) **Traffic Characteristics of Space Science Communications** – Because the application domain involves solar system exploration, the science data being returned from surface assets have traffic characteristics that are atypical of conventional Earth-based wireless networks.

1) Missions will typically generate as much data as the available communication capacity. Each surface asset will generally have enough data to feed continuously to a communications link for as long as it is available.

2) Latency requirements are typically more relaxed and more related to memory capacity than to control loops. In particular, control by human users inherently involves large delay feedback, since Earth is far away and hence the propagation delays are large.

3) There is a very stringent requirement for the reliable transfer of data. Any opportunity for science data return that is lost is costly and the impact of corrupted command data on the forward link may threaten the health and safety of the surface asset.

4) Generally, there is a much greater volume of data returned to Earth than there is sent from Earth, and this is reflected in the Proximity link traffic as well. Proximity-1 is designed to accommodate this anticipated traffic asymmetry. Asymmetric link traffic can be accommodated on Proximity-1 by using different forward and return data rates and frame lengths.

d) **Asymmetry of Communications Resources (Receiver-Driven Operation)** – Relay architectures have proven their cost-effectiveness, greatly reducing the communications resource needs of surface assets. Furthermore, a single relay node may provide service to many such surface assets. This is particularly true of an orbiter, which covers large areas, albeit with episodic links, and has greater visibility with Earth. This relay architecture naturally leads to designs where an orbiting asset will have significantly more communication resources and sophistication than the surface elements. Because the orbiter has greater capabilities, it is natural for it to be the link controller. However, generally the primary traffic flow is telemetry from the landed asset as the source. Therefore this is unconventional in that the 'receiver' of this primary flow (i.e., the orbiter) typically initiates communications.

e) **Large Variety of User Missions** – Because of the expense of space missions, it is advantageous to leverage the relay capability over as many user assets as possible. These assets will be diverse, with a wide range of resources inherently available for communications (including radio, processing and storage capabilities). Examples span from a simple probe to a sophisticated rover or robotic outpost base station. The relay orbiter and its protocol suite must have the ability to accommodate the diversity amongst these heterogeneous user assets.

f) **Efficient Use of Resources** – Placing resources in space is expensive. In particular, energy reserves will be precious for space assets, especially small probes ('scouts') that are particularly dependent on relay support. For this class of relay user, Proximity-1 provides simplex or half-duplex operations in order to minimize energy consumption. Similarly, (buffer) storage and processing resources required for communications should be utilized efficiently. The possibility of sharing the computational and storage resources of the orbiter by the user asset is enabled by the exchange of Proximity-1 messages between asset and orbiter on-board applications.

g) **Proper Balance of Manual versus Autonomous Operations** – Some degree of autonomous operation is necessary for space-based systems since manual interaction is made difficult by the larger propagation delays between Earth and space. While

adaptivity is an essential element of virtually all such protocols, the inability for rapid manual intervention can be particularly constraining for space applications. Moreover, it is desirable to maintain careful deterministic control of the mission. Predictable events, such as epochs where lander-orbiter communication is possible, are programmed, and this 'application' drives the connection management process. On the other hand, occasional transient errors in communications will occur randomly, and should be recovered from in-situ to avoid inefficiencies of using the longer delay DTE links for such purposes. Experience has proven that unpredictable fault conditions may occur after launch, but that clever 'workarounds' can be developed to enable at least partial functionality through remotely commanded changes. However, substantial flexibility is needed in the communication system design. For this reason, Proximity-1 allows the user to control and configure the communication system across many interfaces of the various sublayers of the Proximity-1 Data Link layer. This adds complexity in comparison to conventional Earth-based systems, but is warranted.

While there are additional aspects of space exploration that impact communications system designs, the aspects identified above have resulted in the unique features of the Proximity-1 protocol. These are elaborated in the following sections of this document.

To simplify the discussion, the layer that interfaces above Proximity-1 is referred to as the 'user'. The user will generally be a process in a remote spacecraft that resides within the local vehicle controller or spacecraft Command and Data Handling (C&DH) system (refer to figure 2-2).

Real Mission Implementations

The first few missions to utilize the Proxmity-1 protocol have operations scenarios centered around Mars and include NASA Mars Odyssey Orbiter, BNSC Beagle2 lander, ESA Mars Express Orbiter, NASA Mars Exploration Rovers (MER), and the NASA Mars Reconnaissance Orbiter (MRO). Proximity-1 has generally been accepted by the international space community as the in-situ communications protocol for future Mars missions and a variety of other satellite and planetary missions.

To ensure interoperability and compatibility between mission implementation of the protocol, implementation profiles will be kept. This means that if, for example, designers wish to achieve Proximity-1 interoperability with NASA's Mars Odyssey or ESA's Mars Express, they should look up the implementation profile to which those spacecraft adhere. This is necessary because a core set of features and functions is not declared as mandatory within the Proximity-1 Recommended Standards.

An implementation profile for the Mars Odyssey mission is contained in reference [8], annex C. This profile also applies to the Beagle2, Mars Express, MER and MRO missions. Should significant differences in implementation for future missions arise, new implementation profiles will be created for reference purposes.

2.3 FEATURES OF THE PROXMITY-1 PROTOCOL

2.3.1 GENERAL

The following subsection provides an overview of the features of Proximity-1; they are divided into two categories: (1) session control and (2) data transmission.

2.3.2 PROXIMITY-1 LAYER STRUCTURE

The protocol comprises two layers:

– Physical layer;

– Data Link layer, which is subdivided into the four sublayers (see figure 2-2) as follows:

- • Frame;

- • Data Services;

- • Input/Output;

- • Medium Access Control (MAC).
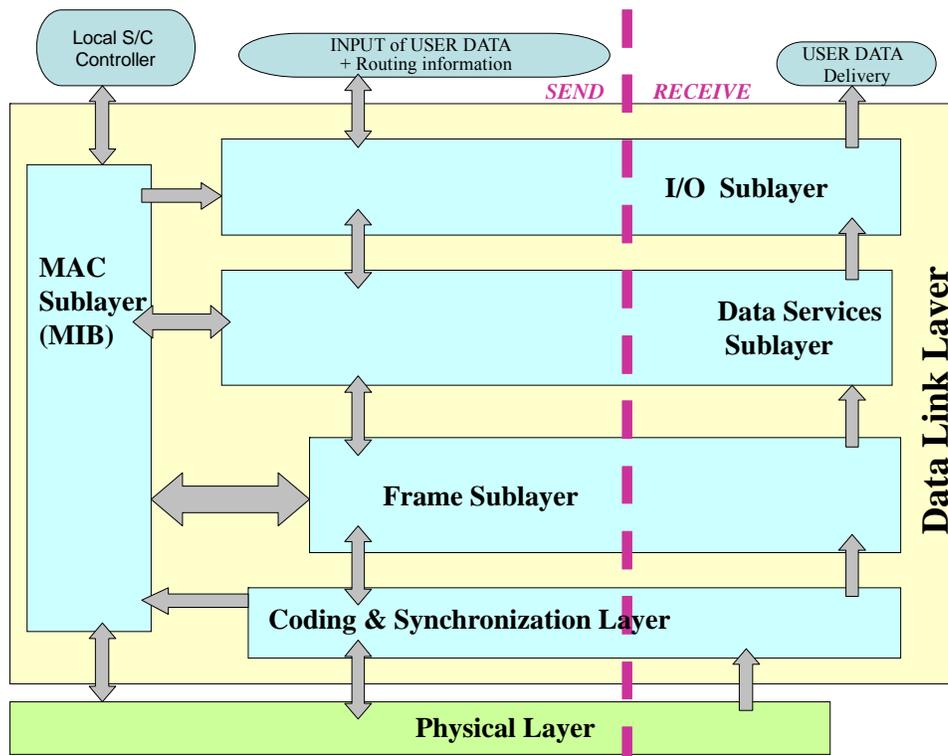
Figure 2-2: Simplified Diagram of Proximity-1 Layers

The formal definition of the whole protocol is spread over three documents referenced in section 1 of this document as follows:

– Physical Layer (reference [10]);

– Coding and Synchronization Sublayer (reference [9]);

– Data Link Layer (reference [8]).

This document supports those formal definition documents by providing a discussion of the essential features and characteristics of the protocol and their application to operational scenarios.

This section of the document addresses the question of what functions/services are provided by Proximity-1 (as distinct from the many different existing and emerging wireless Data Link/Physical-layer protocols) and why they are required in the intended application for Proximity-1.

2.3.3 QUICK LOOK SUMMARY OF FEATURES

The following table is intended to provide an at-a-glance indication of the major components and features of the Proxmity-1 protocol. Further details and descriptions are provided in subsequent sections of this document and the others associated with this protocol as listed in section 1.

Method of achieving/ensuring inter-operability:	Adherence to declared implementation profiles as managed and coordinated by the implementing agency.
	Please note that not all features outlined in the Proximity-1 Protocol Recommended Standards are regarded as mandatory, e.g., the use of coding schemes, etc. It is the responsibility of the implementer to ensure that the appropriate options and features are incorporated in the design to ensure interoperability with other implementations.
	Data Link Layer
Qualities of Service:	Sequence Controlled (using go-back-n, ARQ) (reliable) Expedited (unreliable)
Services:	Data Transfer Time Transfer (for ranging, correlation, and synchronization purposes)

Modes: Full Duplex

Half Duplex

Simplex

Addressing: Spacecraft Identifier (SCID)

Physical Channel ID (PCID)

Port ID

Data Formats: Packets

- CCSDS Space Packets
- SCPS-Network Datagrams
- CCSDS Encapsulation Packets (e.g., IPv6 Datagrams)
- IPv4 Datagrams
- User-defined data units

Frame Length Data Field: Variable up to 2043 octets (not including ASM, Frame Header, or CRC)

Segmentation: Protocol supports segmentation and reassembly of Packets > 2043 octets

Session Characteristics: Link and Session Establishment by Hail process (Duplex only)

Link and Session Termination by a) command directive, b) no more data to send indication or c) loss of signal

Forward (command) Link Data Rates: 1, 2, 4, 8, 16, 32, 64, 128, 256, 1024, 2048, 4096[1] kbps (powers of two)

Return (Telemetry) Link Data Rates: 1, 2, 4, 8, 16, 32, 64, 128, 256, 1024, 2048, 4096[1] kbps (powers of two)

---

[1] Optional – NASA Electra Transceiver capability.

Physical Layer

Modulation Options: FSK[2]

PSK (Coherent/non-coherent)

– PCM data Manchester bi-phase modulation and modulated directly onto the carrier
– Modulation index: $60º \pm 5\%$
– PCM bi-phase-L waveforms mark-to-space ratio: 0.98-1.02

Coding:
– None
– Convolutional (optional), (7, ½) G2 vector inverted
– Concatenated (Reed-Solomon (204,188) with CC (7,½)[3]

Forward Link Frequencies: 435MHz to 450MHz

Return Link Frequencies: 390MHz to 450MHz

Error Rates: Bit Error Rate (BER) = 1×10^{-6} for asynchronous links, e.g., links that do not use the R-S Code

Symbol Error Rate (SER) = 1×10^{-3} (for fixed-length frame applications where R-S coding is performed in the Data Link layer)

The following subsections describe the functional aspects of Proximity-1 from these standpoints:

– Link Session Control;

– Data Transmission;

– Ranging and Timing Services;

– Messaging.

---

[2] Intended for missions such as microprobes with a descoped receiver.
[3] Optional – NASA Electra Transceiver capability.

2.3.4   LINK SESSION CONTROL: ESTABLISHMENT, RE-CONFIGURATION, AND TERMINATION

2.3.4.1   Concept of a Link Session

2.3.4.1.1   General

The Proximity-1 protocol forms a combined link- and Physical layer session between two entities, A and B. The identities of these entities are defined for each session. For full- or half-duplex operation, two-way communication occurs during each session. An important feature of Proximity-1 is that the characteristics (selected parameters) of the two different directions of communications (A to B versus B to A) can be chosen to be very different. These include parameters such as:

–   modulation type;

–   data rate;

–   whether channel encoding (uncoded or convolutional) is used.

These may be different even for half-duplex operation (i.e., when each endpoint uses the same RF center frequency). In addition, Physical and Data Link layer parameters may be changed dynamically during a session.

2.3.4.1.2   Connection-Oriented Data Link Layer

The Proximity-1 protocol is built on the concept of a link session. The protocol is connection-oriented and point-to-point and creates a single physical channel between two communication partners. The time scale of a link session is typically on the order of the duration of time that the orbiter and surface asset are in geometric view of one another. Data exchanges may occur continuously during this data link session. If multiple surface assets are simultaneously in view of the link initiator, and transceiver resources must be shared, then separate sessions will generally be made for each surface asset. This protocol differs from conventional systems, where latency/fairness quality-of-service needs typically call for rapid intermixing of different source Data Link layer PDUs in the course of sharing the radio medium, an approach that can significantly reduce the capacity of the link.

Each Proximity-1 session can be operated in a demand or negotiated communication style. A demand request to establish, change configuration, and terminate the link can be made by either communication partner. Link establishment, data transfer operations, and termination can also occur by negotiation. For link establishment, status reports can be exchanged between nodes to determine communication capabilities. For link termination, each node can inform the other when it has no more data to send ('remote no more data' field in the SET CONTROL PARAMETERS directive) in order to terminate the link at a known data transmission state. Currently a demand hail is used to configure the partnered transceiver's receive and transmit sides.

2.3.4.1.3 Circuit-Switched Physical Layer

Many conventional, particularly wireless, data link protocols are designed to share the physical medium among many users, and operate such that the act of offering a data link SDU implicitly also triggers the actions needed to activate the Physical layer. Correspondingly, the Physical layer activation persists only as long as needed to attempt transmission of the PDU(s) comprising the data link SDU.

Proximity-1 instead operates with a circuit-switched Physical layer, so that physical activation (connection) precedes the ability to convey data link SDU, and is not done automatically: it must be independently initiated. Also the Physical layer actions persist by continuously maintaining IDLE data whether or not data link SDUs are offered; this is the idea behind a '*synchronous channel*'. The rationale for using a circuit switched Physical layer includes:

a) Proximity-1 is point-to-point on the return (telemetry) link; data from several landed assets in the field of view of the orbiter can be acquired by time sharing the return link (assuming single channel receive capability on the orbiter) and dedicating specific intervals of receive time to each asset. There is no time sharing of the channel.

b) Proximity-1 can support one-to-many communication on the forward (command) link. An orbiter can transmit command data to multiple landed assets simultaneously, and each landed asset can be individually addressed based upon its spacecraft ID.

c) Reduction of synchronization overhead between successive transmissions of data link SDUs.

d) For full-duplex operation, Physical layer carrier detection may serve as a clear-to-send (CTS) signal, so that the surface asset (sender) is assured of the receiver's presence.

e) There are typically enough data to fill the transmission channel continuously throughout the session; thus the need to use fill frames or idle bits to maintain synchronization is infrequent.

In fact, real-time interaction between applications in different spacecraft is supported by Data Link layer messaging in Proximity-1. Messages intended for consumption on-board the partnered spacecraft is identified by using a dedicated Port ID. Proximity-1 supports full-duplex communication, although the half-duplex case is also supported. Of course, half-duplex operation alters the capability of maintaining continuous Physical layer synchronization. In this case, the next best possibility is achieved, wherein as long as the communications partner holds access to the channel, continuous Physical layer modulation is maintained. Reacquisition of the channel is needed upon each link turnaround.

2.3.4.2 Session Establishment: The Hailing Process

The Hailing Sequence

In a Proximity-1 communications session there are two parties involved in the process of establishing a communications link:

- the **Caller** who issues a Hail directive (e.g., often an Orbiter);

- the **Responder** who responds to the Caller's directive (e.g., often an asset).

The connection initiation service that Proximity-1 provides to users is known as Hailing. During Hailing the Caller's Local Spacecraft Controller (i.e., user) issues an 'initiate link' directive to the Proximity-1 protocol, which in turn initiates hailing procedures to establish the link.
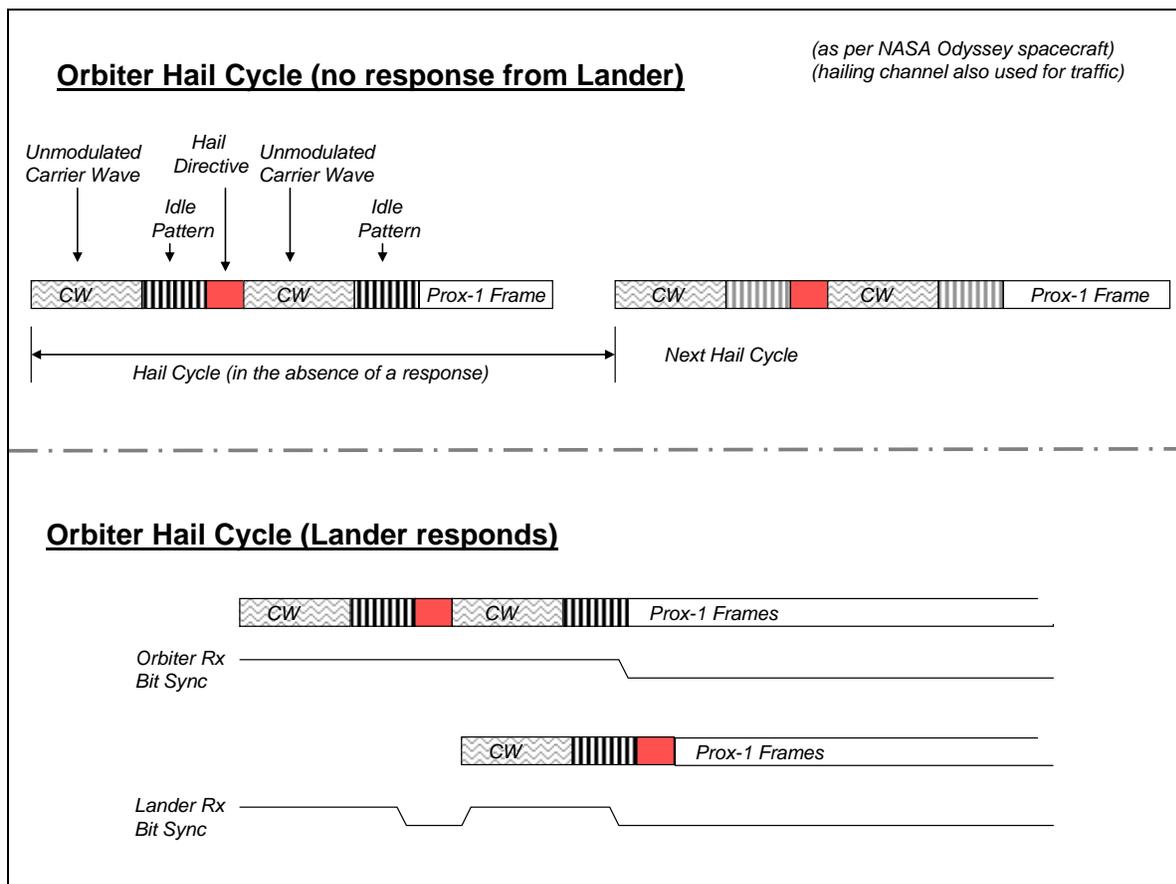

Figure 2-3: Example Hail Sequence Cycle

On receiving an 'initiate link' directive, the Caller's Proximity-1 MAC sublayer issues a hailing signal on a predetermined hailing frequency channel to the Responder. It is the responsibility of the caller to use the correctly pre-determined coding, modulation, and data rate in this process.

The Hail Directive contains at a minimum:

– the source or destination spacecraft ID;

– SET TRANSMITTER PARAMETERS directive;

– SET RECEIVER PARAMETERS directive.

The two directives specify the working forward and return frequencies, coding, modulation, and data rate for the session. It is also assumed that the intended recipient spacecraft (the Responder) must be in listening mode with the appropriate configured default physical parameters ready to receive the hail.

Once the Responder end of the link receives the hail, it responds to the Caller by:

– identifying itself to the caller;

– switching to the working frequency in the hailing message, if it is different from the hailing channel.

After both the caller and the responder established bit-level synchronization (i.e., a Physical layer circuit), the data services phase of the link session, i.e., the transmission of user data, begins.

NOTES

1   Proximity-1 medium access control operates with connection periods that can last for minutes or hours (unlike most medium access protocols that deal with operations at the time-slot level, seconds or milliseconds). This approach allows the spacecraft to conserve energy until communications are known to be possible as a result of geometric relationships dynamically driven by orbital mechanics. When there are relatively few surface assets, the frequency of such opportunities makes this process appropriately performed in a fashion scheduled external to Proximity-1 itself.

2   If a spacecraft is limited to a single transmit/receive frequency pair, then the hailing channel and the working channel will be the same. However, if the transceiver is frequency agile, then the Recommended Standard asserts that hailing is carried out on the hailing channel and immediately after successfully hailing, the users switches to the working channel; the hailing procedure can be carried out in full-duplex mode or half-duplex mode. See 4.1.2 in this document for a detailed description of the handshaking procedure in each case.

3   Using separate hailing and working channels allows more than one communications link to exist at the same time. For purposes of inter-agency and inter-mission interoperability, Proximity-1 defines a specific hailing channel. (See reference [10], subsection 3.4.1.3.2.) Once connected, switching to a separate working channel allows another two entities to establish another connection (and so on). Even though they may be spatially overlapping, interference between simultaneous communications links is avoided.

4   Full-duplex channels may be established having widely different data rates in the opposite direction of data transfer, matched to the asymmetric needs of the traffic flows. Half-duplex operation may also be used, and Proximity-1 allows each spacecraft to use explicit turnaround control (using 'End of Data' indicators) or set different transmit/receive timeout times to accommodate asymmetric traffic flows.

2.3.4.3   Dynamic Run-Time Configuration

During the data service phase of a link session, either user (Caller or Responder) can initiate reconfiguration to ensure the link is operating optimally. Data rates, coding, frequency band, and modulation scheme can all be reconfigured by issuing directives to the local and remote Proximity-1 configuration control entities.

The need for run-time Physical layer reconfiguration can occur for the following reasons:

–   channel re-assignment is taking place to accommodate other Proximity-1 sessions;

–   channel conditions may have changed significantly during the session.

In the latter case, changes in signal-to-noise ratio during a session may make it desirable to change the data rate. Changing to a lower data rate when signal conditions are deteriorating may lead to a longer session. Conversely, changing to a higher data rate when link conditions are improving will increase data throughput.

A user may decide to initiate a reconfiguration process by issuing directives to local and remote Proximity-1 configuration control entities. After receiving the directives, Proximity-1 will suspend all processes and services supporting the transmission of upper-layer SDUs until the appropriate responses are observed. Such responses include:

–   the loss of bit synchronization, as the receiver of the directive adopts the new channel, modulation, and/or data rate;

–   the corresponding changes made by the link initiator needed to re-establish the synchronous channel.

On completion of these changes, data services can resume. Data Link layer protocol parameters, such as go-back-*n* window size, PLCW repeat interval, Receive Duration Timer (half-duplex transmit/receive turnaround timer), etc., are typically pre-configured by both the initiator and the responder prior to session establishment.

Common to both the hailing and the run-time reconfiguration processes is the concept of a 'persistence' activity, which is a handshaking process that requires each action to be followed by the timely observation of a set of expected responses before the next action is taken. Proximity-1 monitors persistent activities and takes the appropriate action if persistence fails (see 3.3.5.5 in this document).

2.3.4.4   Session Termination

Session termination can occur in one of two ways:

a) actively by setting mode to *inactive* and aborting the session, or passively by expiration of the carrier loss timer;

b) by sending notification of 'no more data' to the other node.

Either node involved in a Proximity-1 session may terminate the session by setting Proximity-1 in the *inactive* mode without explicit notification to the other node. In effect, this turns the protocol off for the session. The other node will, as a natural result of the other's action, lose carrier signal on the physical channel and time out, a condition that will terminate the session by default. For full-duplex operation, the carrier loss timer controls how long a spacecraft will wait before terminating the session; for half duplex, both the turnaround timer and the carrier loss timer control the wait time. The shortcoming of such abrupt termination is that there can be ambiguities between nodes as to what was the last piece of data that was reliably transmitted and received.

A graceful method of termination, however, is by reception of a local (LNMD) directive or a remote (SET CONTROL PARAMETERS) 'no more data' directive. Thus, either the local or the remote vehicle controller can initiate the termination process.

In either case, when a session is terminated, the local vehicle controller will be notified and, when Sequence Controlled service is used, be provided with information on which SDU(s) were received and acknowledged by the other node. If insufficient time is allocated at the end of the session to receive the needed acknowledgements to determine the status of all transmitted frames, follow-on Proximity-1 sessions may result in possible duplication or loss of SDU(s). Elimination of this problem is left to the controlling data system.

2.3.5   DATA TRANSMISSION

2.3.5.1   General

Proximity-1 provides message-oriented services. That is, variable-sized SDUs are provided to Proximity-1 by the user, and the SDUs are transported such that the SDU boundaries are recognized at the receiving end. This may be contrasted with a byte-oriented service (e.g., TCP), where a byte queue is identified at the sender, and bytes that the user places in the queue are transported over the link as they are available, but no structure above the byte level is recognized.

2.3.5.2   Proximity-1 Data Transmission Formats

The Proximity Link Transmission Unit (PLTU) is the data-carrying entity used to pass data across the Proximity Link. Its composition is as shown in figure 2-4. It comprises an Attached Synchronization Marker (ASM) followed by a Code Block containing a Version-3

Transfer Frame and a Cyclic Redundancy Check value. The Transfer Frame Header contains information used in the control of the link and its services and also contains data about the format and nature of the data contained in the Data Field. Details of the PLTU components, their function and form can be found in the Recommended Standards as listed in section 1.

The reader's attention is drawn to the Data Field Construction ID (DFC ID) which indicates the form of the user data (not protocol data) held in the Data Field of the Transfer Frame. Figure 2-4 lists the four forms. There are varying constraints and conditions associated with each of these data types more details of which can be found in the Recommended Standards:

- **Service Data Units** must comprise an integer number of unsegmented packets.

- **Segmented data** must be preceded by header information indicating segment type:

  - first segment;

  - continuing segment;

  - last segment;

  - unsegmented.

- The **Reserved** type is not to be used until it has been defined for use by CCSDS.

- The **User Defined Data** type has no constraints on the format of the data it contains. By implication the data formatting will be known and decoded at the user's applications level above the protocol.
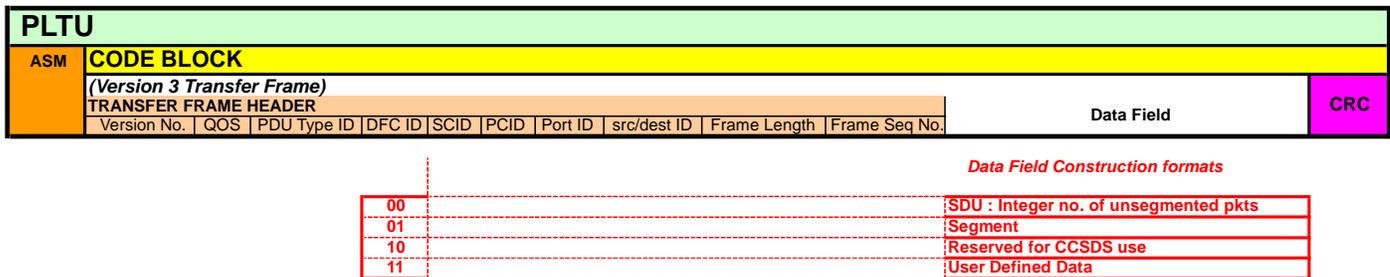
Figure 2-4: Proximity-1 PLTU Format

Further discussion and information on the nature and handling of SDUs can also be found in 2.3.5.2.

Figure 2-5: Proximity-1 Data Formats

2.3.5.3   Use of Synchronous Channel and Asynchronous Data Link

Once a link has been established, the RF transmission is dedicated to the communicating entities (users) until explicit termination or ungraceful closure. The Data Link layer is connection-oriented and the Physical layer operates in circuit switched fashion; bit-level synchronization is maintained continuously during the session except during very brief periods within a communications change. In Proximity-1 this is what is meant by a

synchronous physical channel. After link establishment has occurred, an asynchronous data link is run over this synchronous physical channel. When data is available for transmission, it is carried over the data link in variable-length frames within PLTUs. The use of variable-length frames simplifies the transmission of variable-length SDUs. When there is no data to transmit, IDLE data are sent across the link to keep the physical channel occupied, thereby maintaining the synchronous physical channel.

2.3.5.4   Addressing: Spacecraft ID, Port ID, and Physical Channels

Proximity-1 uses three addressing labels:

  – Spacecraft ID;

  – Physical Channel ID;

  – Port ID.

The Spacecraft ID field identifies either the source or destination ID of a Proximity-1 PDU by using a source/destination flag and an identifier. This flag allows Proximity-1 to designate the 'destination' of a command message, as in Telecommand, or designate the 'source' from which data came, as in Telemetry. In the prototypical 'orbiter/landed asset' scenario, the destination ID allows multiple landers to discriminate their content on a shared forward link, while the source ID allows the orbiter to distinguish data received from multiple landers.

The Port ID provides the means to route user data internally (at the transceiver's baseband output interface) to specific logical ports, such as applications or transport processes, or to physical ports, such as on-board buses or physical connections (including hardware command decoders). Well-defined Port IDs for both the forward and return link are assigned to specific data types in the Space Link Identifiers Recommended Standard (reference [14]).

Proximity-1 provides two independently multiplexed physical channels, each capable of supporting both the Sequence Controlled and Expedited services. In addition, Proximity-1 supports requests for PLCWs based upon the physical channel in use. However, operations involving the simultaneous use of both physical channels is not part of the Recommended Standard.

2.3.5.5   Proximity-1 Service Data Unit

Once data service begins, the protocol can accept three kinds of SDUs from the upper layer:

  a) Packets, i.e., CCSDS Space Packets, SCPS-NP datagrams, IPv4 packets, CCSDS Encapsulation packets(includes IPv6);

  b) user-defined data units of unknown format; and

  c) Proximity-1 protocol directives pre-formatted by the user as SPDUs.

The first two categories are considered to be *user data* by Proximity-1, and the third is considered to be *supervisory data* related to Proximity-1 operation.

Proximity-1 provides the ability to fit several SDUs into a single frame, spanning SDUs across multiple frames, or segmenting a single SDU into multiple frames. There are conditions that govern whether two SDUs can share a single frame.

If the SDU contains a packet, the upper-layer application that receives this SDU can determine which protocol to route this SDU to by evaluating the version ID (which specifies the protocol) within the packet header. This functionality is important when multiple upper-layer protocols are sharing the same data link protocol, i.e., Proximity-1.

Supervisory Protocol Data Units

SPDUs are SDUs that carry data related to the administration and operation of the Proximity Link by the protocol. Figure 2-5 shows how SPDUs fit into the Proximity-1 Transfer Frame schema. From this diagram it can be seen that SPDUs can be of variable or fixed length, always use the Expedited Quality of Service (QOS), and come in a variety of types including:

− Directives/Reports/PLCWs;

− Time Distribution PDUs;

− Status Reports.

The SPDU Format ID indicates whether it is of variable or fixed length and the SPDU Type ID field indicates into which of the three types listed above the data fit. Note that with the Directives/Reports/PLCWs, the last three bits of the data field must be decoded to determine the nature of the data field contents and their format. The last three bits were chosen because of the heritage design of the directive set within the CE 505A transceiver used on 2001 NASA/JPL Mars Odyssey Orbiter.

(See reference [8], subsections 3.2.4 through 3.2.8, on data field construction rules for packets, segments, user-defined data, and SPDUs.)

2.3.6   TWO GRADES OF LINK-LAYER SERVICE

2.3.6.1   General

On any given physical channel, Proximity-1 can provide two grades of services: (1) Expedited and (2) Sequence Controlled. Both Expedited and Sequence Controlled frames can be multiplexed onto the same physical channel. User data can be transported using either grade of service, while protocol directives always use the Expedited service.

NOTE – In the following subsections, the relationship between SDUs and PDUs is as per the OSI reference model (as expressed below to aid understanding).

To enable SDUs to be passed from one layer to the next (lower) layer, Protocol Control Information (PCI) is added as a header thereby forming a PDU (a process known as 'encapsulation' or 'layering').
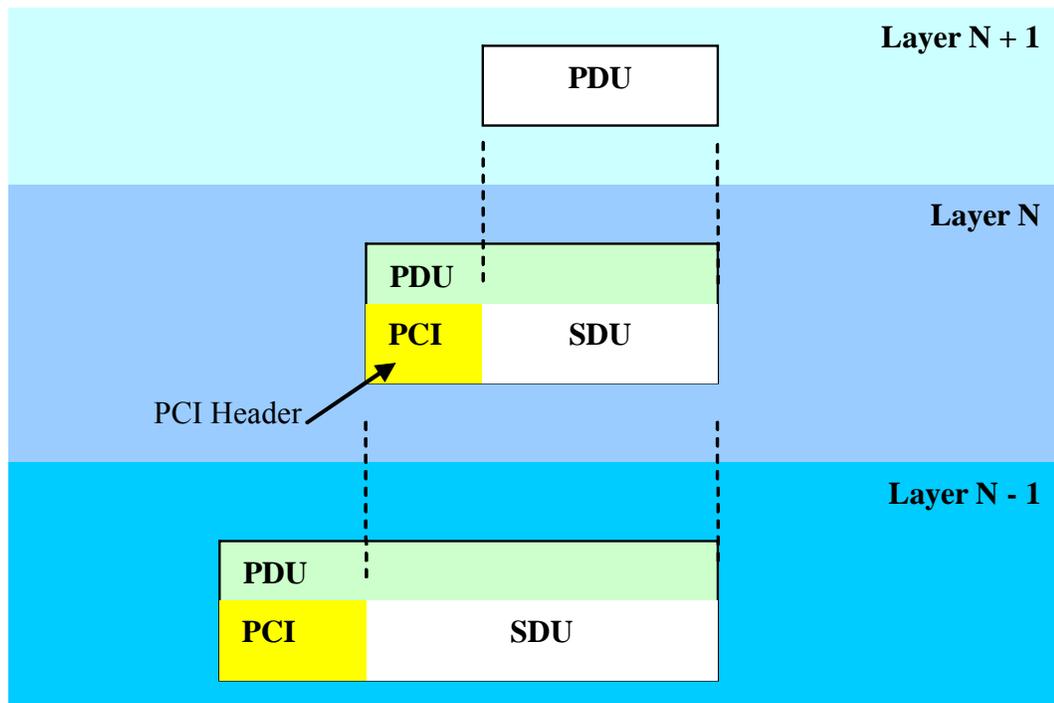
Figure 2-6: Relationship between PDUs and SDUs

Thus, an SDU is coupled with a PCI header to become a PDU. When that PDU is passed to the next layer down it becomes an SDU at that layer. The Protocol Control Information informs that layer in which it resides what processes and values must be applied to the SDU at the corresponding layer at the receiving end of the link.

2.3.6.2 Expedited Services

Expedited service delivers SDUs in the order they are received from the upper layer because Proximity-1 operates as a point-to-point link and delivers each SDU in a first-in-first-out (FIFO) fashion. The delivery mechanism converts each SDU (packet or user-defined data) into PDU(s) (Proximity-1 frames) and applies error detection on each PDU to ensure data integrity. A PDU will not be accepted if any error is detected. Since a 32-bit CRC is used to detect errors in the frame, the undetected bit error rate is approximately 10^{-11}. (See annex A of reference [9].) Thus with high probability, any SDU extracted from accepted PDUs will be error-free. However, since dropped PDUs will not be retransmitted, SDU losses may occur.

It is worthwhile to note that Proximity-1 will deliver only *complete* SDUs to the user; if an Expedited SDU is broken into multiple PDUs, and any one of these PDUs is lost, then the entire SDU is not delivered. *Expedited service does not deliver partial SDUs with errors or 'holes'*. Therefore, if partial delivery of an image is preferred to image loss, it is better to submit the image data to Proximity-1 as a collection of smaller SDUs, rather than one large SDU.

Expedited service maintains separate sequence numbers for each SDU stream based on the Physical Channel ID. All Expedited SDUs (and associated PDUs) receive priority transmission over all Sequence Controlled SDUs/PDUs for a given physical channel.

2.3.6.3 Sequence Controlled Service

The Sequence Controlled service may be selected for SDUs that carry user data. Unlike the Expedited service, retransmission of PDUs is possible using a go-back-*n* Automatic Repeat Queuing (ARQ) algorithm. Proximity-1 thereby guarantees loss-free, acknowledged, duplicate-free, and high-integrity delivery of SDU(s) in the order they are received *within* each Proximity-1 session, providing there are no COP-P resynchronizations. Expressed simply, data are reliably transferred across the space link and delivered in order, without gaps, errors, or duplications. However, if a COP-P resynchronization has occurred, missing or duplicate data may occur during the session because of factors outside of the control of the protocol.

In contrast to a selective ARQ algorithm, the use of the relatively simple go-back-*n* mechanism is appropriate since efficiencies are not affected in the relatively short propagation delay links intended for Proximity-1. The complexity of implementation is also minimized, which is important for resource-limited assets.

PDUs are queued separately according to their grade of service and Expedited SDUs (and associated PDUs) receive priority transmission over all Sequence Controlled SDUs/PDUs for a given physical channel. This means that Expedited service PDUs may arrive earlier than Sequence Controlled SDUs even if they are submitted later.

2.3.7 DELIVERY ACKNOWLEDGEMENT TO PROXIMITY-1 USER

The Proximity-1 Recommended Standards require that the sending user must be informed of the status of each SDU delivery. For Expedited service, when all PDUs associated with an SDU have been transmitted, the sending user is notified. For Sequence Controlled service, the sending user is notified when all frames associated with the SDU have been correctly received and acknowledged. This feature is essential to the seamless data delivery across multiple Proximity-1 link sessions. It is important to note that duplicate-free and/or gapless delivery cannot be guaranteed across multiple Proximity-1 link sessions in the event that any link session is terminated ungracefully. In this case there can be ambiguity about the delivery status of outstanding frames and the associated SDU(s).

Note that the delivery acknowledgement feature of Proximity-1 requires that the PDU (frame) sequence numbers be associated with the SDU from which they were created. For more details, see discussion in 4.1.6.6 in this document.

2.3.8 RANGING AND TIMING SERVICES

Proximity-1 also provides ranging and timing services to users. By time stamping the departure and arrival time of Proximity-1 transfer frames exchanged between the two spacecraft, during a commanded interval, the round-trip time between two Proximity spacecraft can be derived accurately.

The accuracy of such calculation depends on

- the stability of the clock on each spacecraft;

- the radio's ability to determine accurately the time of the trailing edge of a designated bit within the frame used for time tagging;

- the ability to determine accurately the delays incurred during transmit/receive processing of the transfer frames that are time tagged.

Proximity-1 provides both a mechanism for Proximity time correlation as well as exchanging time between spacecraft. Measuring the round-trip time also provides a means to estimate range between spacecraft based on the propagation delay of the speed of light.

2.4   MULTI-CHANNEL AND MULTI-CONNECTION EXTENSIONS

Proximity-1 provides a very efficient means for an orbiter to act as a relay for a wide-range of different assets. The protocol is designed to provide a single link session at a time between one or more spacecraft on the forward link and only one spacecraft on the return link when a single physical channel is available. Multiple transmitters/receivers are required on the return link to enable one entity to communicate simultaneously with two or more other entities. However, Proximity-1 will allow multiple links to operate simultaneously even though there is spatial overlap, by providing the means to establish separate, non-interfering 'working channels'.

Also, if an entity needs to communicate with two other entities that simultaneously fall within its RF footprint, it may do so in a time sequenced fashion. This is done by establishing a link session with one, performing the communications, closing that session, and then establishing a second session with the second entity to perform the remaining communications. The entity that is not part of the individual session will either ignore hails and frames received not having the destination SCID, or can be turned off by a higher-layer scheduler knowledgeable of the sequential nature of the procedure.

It is conceivable to design a relay orbiter with resources that would allow multiple links to operate concurrently. Extensions to the protocol toward achieving this goal are discussed in 5.3.3.

More generally, there is no requirement that the relay element be a satellite; it may be a base station that communicates with a collection of mobile or fixed assets in a star topology. In fact, as noted earlier, the master controller need not provide any relay function, i.e., forwarding the data to another entity after the Proximity-1 session has ended; it may simply be engaged in a two-way conversation. Many other applications scenarios may be considered, not limited to space environments. The flexibility of the interface to the user offers the potential to extend the baseline protocol for more sophisticated data link control in its operation.

3 PROTOCOL ARCHITECTURE AND SERVICES

3.1 OVERVIEW

The Proximity-1 protocol specifies the Data Link layer, Coding and Synchronization sublayer, and the Physical layer of the OSI reference model as well as their associated control/management functionalities. The Data Link layer functionalities are implemented by four components, or sublayers:

a) I/O sublayer;

b) Data Services sublayer;

c) Frame sublayer;

d) Link Connection/Configuration Control (called 'MAC' sublayer in the Recommended Standard).
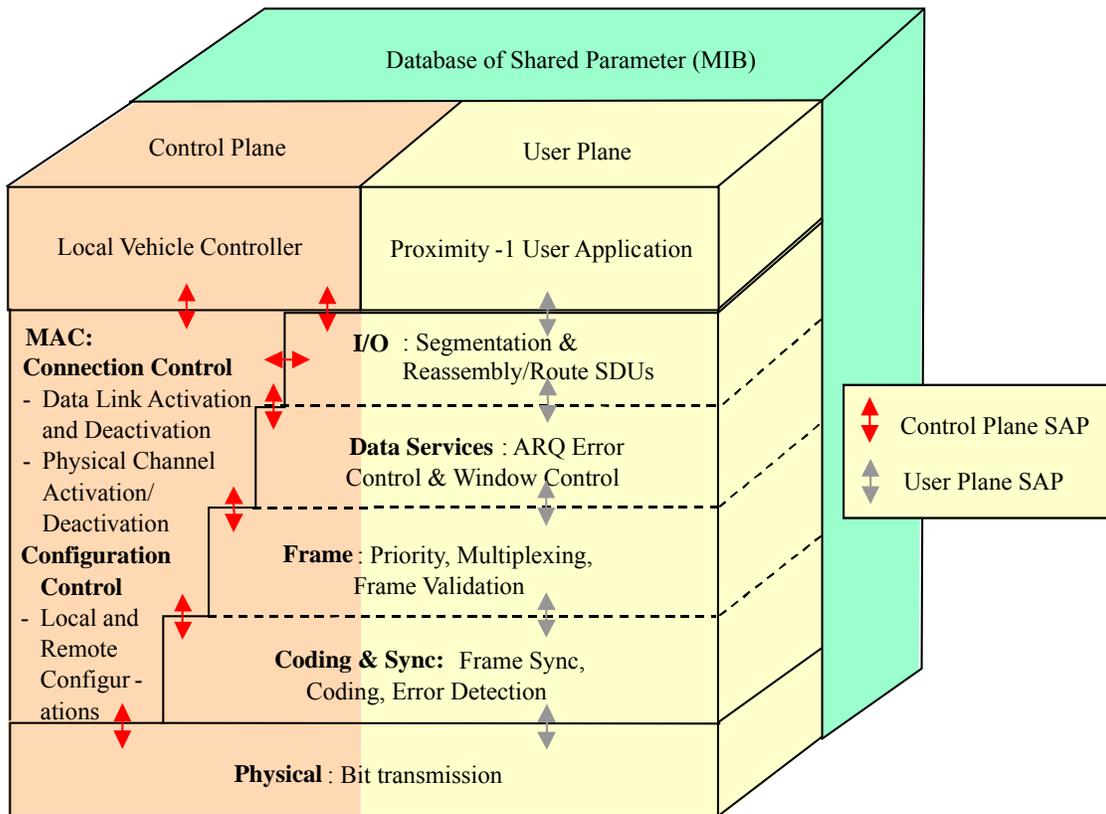
**Figure 3-1: Proximity-1 Architecture and Functional Models**

Figure 3-1 depicts the Proximity-1 protocol layers, sublayers, and their functional relationships in both the user and control planes, and the Services Access Points between the user and the sublayers. The following subsections give a service-based, functional description for each protocol component in the control plane and the user plane.

3.2 PROXIMITY-1 USER AND SERVICES

3.2.1 GENERAL

The Proximity-1 protocol operates at the request of a 'user'. Possible users of Proximity-1 include higher-layer protocols, e.g., CFDP, SCPS-NP, SCPS-SP, IPv4, IPv6, etc. Proximity-1 provides two basic classes of service: management/control and application. Management/control services allow the user to direct and monitor the establishment, configuration, and termination process of Proximity-1 sessions; application services accept data units from the user and transfer them to their designated recipient using different grades of services. If the functional relationship is considered between these two classes of services and the Proximity-1 protocol stack, it is clear that the user plane enables the application services, and the control plane enables the management/control services. For this reason, the management/control services are referred to as 'control plane services', and application services are referred to 'user plane services'.

A unique feature of the Proximity-1 Recommended Standards is the extensiveness of control plane services. Proximity-1 users are given a great deal of control over the internal operation of the protocol. This is accomplished by increasing *inter-layer coupling*, thus making direct, manual control of lower-layer parameters possible. This is typically not the case in a terrestrial network system where strict layer transparency is considered a virtue. The ability of a user explicitly to control lower-layer parameters is desirable for space-based operations, since the potential for manual intervention to rectify unforeseen faults or conditions counter-balances the additional complexity. Therefore the Proximity-1 user, rather than just providing the data and letting the lower layer handle all of the operational complexity, retains control of the Physical layer and the activation, deactivation, and configuration of both the physical and data link connections.

3.2.2 CONTROL PLANE SERVICES

3.2.2.1 General

Control plane services includes the ability for the user to establish, configure, and terminate a session, *regardless of whether there is data to send or not.* The user has the ability to control a remote Proximity-1 entity, request and receive status reports about the on-going session, and request a Proximity-1 time tag interchange.

Figure 3-2 shows conceptually how the internal structure of Proximity-1 looks. The building blocks of all control plane services are a set of *primitives*: control points within Proximity-1 that directly affect the operation of each user plane sublayers. Some of these primitives control points are accessible to the user. Some are in turn controlled by various state variables and their associated sequential logic to perform *basic control functions* for the user. The 'half-duplex transmission control' and 'physical layer comm change' are both examples of basic functions described in the Recommended Standard. The set of accessible primitives and basic control functions make up the core of Proximity-1 control plane services, and they can be invoked locally or remotely by a set of *directives*.

In order to provide the maximum degree of manual control and flexibility, control plane services are not highly automated 'single-button' functions; control procedures, such as hailing and run-time re-configuration, which require sequential or simultaneous execution of multiple local and remote directives may require interaction with the local spacecraft controller. See annex D in reference [8] for the protocol notifications to the vehicle controller.

Figure 3-2: Control Plane Services (Timing/Ranging Not Included)

The following subsections describe two types of control plane services:

a) services invoked by directives:

    1) issued to a local Proximity-1 entity at the user-MAC sublayer interface, or

    2) issued to a remote Proximity-1 entity by sending a **Supervisory PDU, which contains protocol directive(s)** at the user-I/O sublayer interface;

b) status report triggered by Proximity-1 events.

Note that although many SPDUs are issued by the MAC sublayer, they can also be issued by the user manually, thus having an extensive set of control capabilities on Proximity-1 operations. The follow subsections provide a brief summary of the control plane services provided by the Proximity-1 directives and those triggered by Proximity-1 events.

3.2.2.2   Services Invoked by Proximity-1 Directives

a) SET TRANSMITTER PARAMETERS directive

Set transmission mode, data rate, encoding, modulation, and frequency. This directive is used in link establishment, and Physical layer changes (e.g., data rate, frequency) during data services.

b) SET RECEIVER PARAMETERS directive

Set receiver mode, data rate, decoding, modulation, and frequency. This directive is used in link establishment, and Physical layer changes (e.g., data rate, frequency) during data services.

c) SET MODE directive

1) Turn off transceiver (*inactive*).

2) Enable transmitter in asynchronous physical channel mode (bit synchronization not maintained) and receiver (*connecting-T*).

3) Enable receiver only (*connecting-L*).

4) Enable both transmitter in synchronous physical channel and receiver (*active*).

d) SET DUPLEX directive

Set full-duplex, half-duplex, or simplex-receive, simplex-transmit operation.

e) SET CONTROL PARAMETERS directive

1) Request *N* time tags taken.

2) Set link directionality (e.g., full, half duplex, simplex).

3) Inform the remote node that there is no more data available to send (link termination process).

4) Switch transmit/receive role in half-duplex operation by means of token pass.

f) REPORT REQUEST directive

1) Request various status reports – report format to be specified by the user.

2) Time tagging – request time tagging of Proximity frames upon ingress and egress by both transceivers (used for Proximity Time Correlation).

3) Request PLCW – request a PLCW from the partnered node per Physical Channel ID.

g) TIME DISTRIBUTION – this directive category includes several specific directives and services needed for the user to perform timing computations:

1) TIME TRANSFER;

2) BROADCAST SPACECRAFT CLOCK TIME.

There are other types of directives that can be issued by the upper-layer user but are only designed to facilitate protocol interaction between internal Proximity-1 processes. Directives such as the **SET V(R) directive** is intended for use by the COP-P to force resynchronization with its communication partner. The **REPORT REQUEST directive** has options for requesting a status report about the status of data delivery per physical channel, such as sequence numbers. This directive is intended for use by the Data Services sublayer to prompt a remote Proximity-1 entity to generate an acknowledgement for outstanding frames.

3.2.2.3 Report Triggered by Proximity-1 Events

When certain events occur, Proximity-1 will automatically notify the user. The following are some of the events that will trigger automatic notification:

a) success or failure of a persistent activity (hailing, reconfiguration of frequency, data rate, modulation, etc.);

b) termination of session due either to explicit notification or to prolonged loss of carrier signal.

See annex D in reference [8] for a complete list of protocol notifications to the vehicle controller.

3.2.3 USER PLANE SERVICES

3.2.3.1 General

The basic user plane service is the delivery of SDUs *between Proximity-1 users* between two spacecraft. The delivery services can be described and distinguished based on the following four features: (1) qualities of service, (2) number of connections, (3) SDU routing and 'protocol identification', and (4) delivery notification.

3.2.3.2 Qualities of Service

There are two qualities of service for SDUs delivered over the Proximity-1 link which allows for the multiplexing of both Sequence Controlled and Expedited frames:

a) Sequence Controlled – Provides in order, without gaps, without errors, or duplicates delivery of complete SDUs within a session.[4] SDU reliability is achieved by running the go-back-*n* ARQ algorithm, i.e., COP-P on the Proximity-1 PDU (frame).

b) Expedited – Ensures error free delivery of complete SDUs within a Proximity-1 session with the possibility of SDU loss. However, Expedited PDUs (frames) are given higher transmission priority than Sequence Controlled PDUs (frames).

For both qualities of service, *corrupted or incomplete SDUs are not delivered to the user, nor will their loss be directly indicated. However, the loss of CCSDS Space Packets can be discovered by the discontinuity in the packet sequence count.* For any application using the Proximity-1 protocol, the SDU size defines the granularity of data loss. A mismatch between the application's tolerance for data loss and the Proximity-1 SDU size may cause unnecessary data loss when using Expedited service.

For example, if the Proximity-1 user sends an image of a landscape using the entire image as a single SDU, then the unit of loss would be the entire image. The loss of any Proximity-1 PDU(s) (frames) associated with an SDU will cause that entire SDU, therefore the entire image, to be discarded. On the other hand, if the image is divided into multiple SDUs, each containing independently decodable segments of the image, then it is possible to deliver this image with gaps using the Expedited service.

3.2.3.3 Multiplexing of Frames within a Connection

3.2.3.3.1 General

The notion of SDU 'order' is preserved within each physical connection using independently incremented frame sequence numbers for each quality of service. For both Sequence Controlled and Expedited services, the Physical Channel ID specifies the physical connection to which an SDU belongs.

3.2.3.3.2 SDU Routing and Protocol Identification

Proximity-1 delivers each SDU to the receiver-user through a specified output port. The Port ID provides a kind of 'hard-wired' on-board routing functionality so that an SDU can be directed to a particular processor, memory buffer, or storage device. However, besides serving as an on-board routing address, the Port ID field can also be used for protocol identification purposes by uniquely associating a port to an upper-layer protocol.

---

[4] Assuming no resynchronization of the Sequence Controlled service has occurred during the session.

3.2.3.3.3   Delivery Notification to the Data Supplier

The I/O sublayer is required to provide delivery status for each SDU to the data supplier. If an SDU is delivered using Expedited service, then the data supplier will be notified about which Proximity-1 frames containing the SDU were transmitted; there is no guarantee, however, that the SDU will be received successfully. If an SDU is delivered using Sequence Controlled service, then Proximity-1 will keep track of the last frame sequence numbers of a series of frames that span the SDU and notify the user when all frames associated with a particular SDU have been received correctly and positively acknowledged.

This feature is an important part of the session termination process because it tells the data supplier which SDUs have been transmitted and acknowledged. This enables the seamless resumption of SDU transfer from the point of interruption to when the next session begins. For delivery notification to work properly, both ends of the communication link must allow sufficient time at the end of a session to generate and receive the needed feedback information before terminating the connection. Abrupt termination of a session may create ambiguity about the delivery status and therefore cause loss or duplicated delivery of SDUs when the Proximity-1 link is re-established.

3.3   PROXIMITY-1 PROTOCOL LAYERED FUNCTIONALITY

3.3.1   I/O SUBLAYER

3.3.1.1   General

The I/O sublayer performs the following send functions:

a) provides the interface to accept all the information necessary (e.g., Port ID, PCID, QOS) to transfer SDUs from the sender-user;

b) packages an SDU into Proximity-1 transfer frame(s) and labels each frame with the appropriate sequence number;

c) delivers Proximity-1 transfer frames to the lower layers through separate FIFO queues maintained for each quality of service;

d) parses large SDUs into segments compatible with the maximum transfer unit size over the link;

e) provides reports to the sender-user on the status of SDU delivery:

   1) Expedited mode – notify sender-user when transmission of all PLTUs associated with an SDU over the physical channel has been completed;

   2) Sequence Controlled mode – notify sender-user when all frames associated with a *complete* SDU has been received and positively acknowledged by the receiver.

The I/O sublayer performs the following receive functions:

a) receives Proximity-1 transfer frames from lower layers and reconstructs the original SDUs for delivery to the receiver-user through the output Port ID specified by the sender-user;

b) reassembles SDUs that were segmented.

3.3.1.2 Receiving SDUs from the User

The I/O sublayer provides the service access point (SAP) through which the user can supply SDUs to the Proximity-1 Data Link layer. Acceptable SDU formats include CCSDS space packets, CCSDS Encapsulation packet, SCPS-NP packets, IPv4 and IPv6 datagrams, etc., as well as other user-defined formats unknown to the protocol. The I/O sublayer interface also allows the user to specify, along with each SDU, the quality of service (Expedited or Sequence Controlled), the output Port ID through which the SDU is to be delivered on the receive side, the physical channel, the PDU type (whether data or protocol directives are contained) and the U-frame data construction rules by which the SDU is contained within Proximity-1 transfer frame(s).[5]

3.3.1.3 SDU Delivery (Sending Side)

For the sender of an SDU, the I/O sublayer will receive information from the Data Services sublayer about which Expedited frames are sent and which Sequence Controlled frames are both sent and positively acknowledged. The I/O sublayer will translate the report of frames received into their equivalent in terms of SDUs successfully transferred and update the user on the delivery status for each SDU. By keeping track of the frame number that contained the last part of an SDU, the sender-user can track when an SDU is reliably transferred by knowing if that last frame is acknowledged. On the receiving side of the link, U-frames will be delivered from the lower layer (Data Services sublayer) to the I/O sublayer. The I/O sublayer will unwrap the U-frame headers and rebuild the original SDU for delivery to the user through the user-specified Port ID. Only complete SDUs will be delivered to the user. In Proximity-1, P-frames that originated from the sender-user will only be delivered to the MAC sublayer on the receiving side. The I/O sublayer does not process any P-frames on the receiving side.

---

[5] See the data field construction rules in reference [8].

3.3.2 DATA SERVICES SUBLAYER

3.3.2.1 Overview

The Data Services sublayer's primary functions are:

a) Execute the necessary sequential logic and buffer management tasks to deliver Proximity-1 frames using: (1) go-back-n ARQ, i.e., Sequence Controlled QOS and (2) Expedited QOS.

b) Provide information on the sequence numbers of the frames that are transmitted and also those that are acknowledged to the I/O sublayer.

3.3.2.2 COP-P Process

3.3.2.2.1 General

The Data Services sublayer implements the go-back-n algorithm using a COP-P process. The COP-P process consists of a sender process called FOP-P and a receiver process called FARM-P.

3.3.2.2.2 FARM-P

Since the lower layer and sublayers (Physical, Coding and Synchronization, and Frame) will eliminate frames received with an error or invalid format and contents, all frames submitted to the FARM-P process are assumed to be error free. Therefore the responsibility of the FARM-P process is to check whether a frame arrives in or out of order. Frames that arrived in order are accepted; those that arrive out of order are rejected. The FARM-P reports to the sender the value of the next sequenced controlled frame it is expecting, i.e., the sequence number of the latest frame that has been accepted plus one. The report value is sent in the PLCW.

3.3.2.2.3 FOP-P

The FOP-P process is in charge of sending either new frames or previously sent frames based on the feedback from the FARM-P process on the receiver side. This means that frames that are sent but not yet positively acknowledged must be stored in a temporary buffer in case retransmission is necessary, i.e., the Sent queue. The FOP-P process must also throttle back transmissions of frames when the number of outstanding frames reaches the window size (N). For full-duplex operation, the choice of N depends on the round-trip time, the frame sizes and the data rates on the forward and return links; for half-duplex operation, channel error rate plays a more important role because of the additional time required to change sender and receiver rolls.

3.3.3 FRAME SUBLAYER

3.3.3.1 General

The Frame sublayer performs the following functions:

a) Multiplex P/U-frames supplied by the Data Services sublayer, P-frames from the MAC sublayer, and P-frames (PLCW) into a single stream for transmission over the physical channel. The following priority is applied during multiplexing:

　　1) P-frames from the MAC queue in the MAC sublayer;

　　2) a PLCW/status report if the previous frame sent was a U-frame;

　　3) Expedited frames;

　　4) Sequence Controlled frames (first from the Sent_Queue else a new frame);

　　5) a PLCW/status report if the previous frame sent was NOT a U-frame.

　　NOTE – Additional conditions are given in reference [8], subsection 4.1.2.2.

b) Create PLCW (feedback frames that carry FARM-P state information needed for the go-back-n ARQ algorithm) by constructing a separate P-frame.

c) Validate received frames based upon static fields in the transfer frame header.

d) Route U-frames and PLCW SPDUs to the Data Services sublayer; other SPDUs are sent to the MAC sublayer.

3.3.3.2 Asynchronous Data Link

3.3.3.2.1 General

Since Proximity-1 runs over an asynchronous data link, frames are of variable length to match the size of the packets transferred. When these frames are not generated fast enough to fill the physical channel with bits, the Physical layer inserts the idle pattern, defined in the Coding and Synchronization sublayer (reference [9]). The frame size is variable as long as it is smaller than the maximum allowed on the link. There are three defined frame data field construction rules: segmented, unsegmented, and user-defined data (octet-aligned stream).

3.3.3.2.2 Segmented Packets

a) For segmented packets, each frame can contain one segment from one packet; therefore each segment is aligned within a frame.

b) Each frame containing a segmented packet has an identifier ('sequence flag' in the Segment Header) that specifies whether the segment is the first, one of the middle, or the last segment of a packet. Combined with the frame length information, the length

of a packet can be computed, if all of the segments are provided. Therefore, transferring segmented packets is recommended only while using the Sequence Controlled service, since all of the segments are guaranteed to be delivered and Proximity-1 delivers only complete packets to the user. Note that the loss of any intermediate segments would be undetectable in Expedited service unless a check sum were calculated over the entire packet before transmission.

c) A segmented packet is reassembled from frames containing the same routing ID, i.e., the same Physical Channel ID, Port ID, and 'pseudo packet ID' (in the Segment Header).

d) Segments from other packets can be multiplexed into the stream of frames as long as they contain a different PCID or Port ID.

3.3.3.2.3 Unsegmented Packets

When a packet can fit within a single frame, segmentation is unnecessary. Since the frame size is variable, there is also no need to include filler data within the packet. If multiple integral packets can fit within the data field of a single frame, they must share the same Physical Channel ID, Port ID, QOS, PDU type, and DFC ID, all of which are fields within the transfer frame header.

3.3.3.2.4 User-Defined Data

User-defined data is placed into a frame without fill; it is up to the user to provide user-defined SDUs that are no larger than the maximum data field size, if it is intended for the frame and the SDU to be in alignment. If larger, user-defined data can be segmented.

3.3.3.3 Sequence Number Assignment

In Proximity-1, sequence numbers are essential for tracking the order of frame delivery, detecting gaps, and making re-transmission requests. The frame sequence number is an eight-bit field in the frame header; a sequence number is assigned to each frame based on the quality of service requested:

a) Sequence Controlled frame – Distinct sequence numbers (SEQ_CTRL_FSN) are incremented monotonically for each physical channel;

b) Expedited frame – Distinct sequence numbers (EXP_FSN) are incremented monotonically for each physical channel.

Note that both of these sequence numbers increment independently of Port ID. The SEQ_CTRL_FSN enables the Sequence Controlled process to number sequentially and then check the sequence of incoming Sequence Controlled transfer frames. The EXP_FSN is not used in the frame-validation process but is required for correlations associated with Proximity-1 timing services.

3.3.3.4 Delivery of Proximity-1 Frames to Lower Sublayer

Proximity-1 frames, whether they are U-frames or P-frames, are delivered to the lower sublayer via two FIFO queues: (1) Expedited and (2) Sequence Controlled, based on the QOS selected. Frames are queued in the order in which they are generated from the SDUs, and their position in the queues is completely independent of Physical Channel ID and Port ID.

3.3.4 CODING AND SYNCHRONIZATION SUBLAYER

The Coding and Synchronization sublayer has two basic functions:

a) frame delimiting:

1) on the transmitting side, marks the beginning of a frame with an Attached Synchronization Marker (ASM) on transmission; and

2) on the receiving side, searches the bit stream supplied by the Physical layer for the ASM, and retrieves the frame using the value of the frame length in the frame header;

b) error detection:

1) on the transmitting side, computes the CRC and appends it to the end of the Proximity-1 frame; and

2) on the receiving side, applies error detection:

– if one or more errors are detected, the frame is rejected;

– otherwise the frame is extracted and submitted to the Frame sublayer.

This sublayer performs the conversion between a Proximity-1 frame and a Proximity Link Transmission Unit (PLTU), which consists of an ASM, a frame, and a CRC.

3.3.5 MAC SUBLAYER

3.3.5.1 General

The MAC sublayer is quite different from the traditional medium access control in two aspects:

a) The traditional MAC function is connectionless and operates in the user plane; it is concerned only with scheduling individual PDUs for transmissions on a shared physical medium. However, the MAC sublayer in Proximity-1 operates in a connection-oriented fashion, and it is primarily a control/management-plane function concerned with establishing a 'link session'.

b) Most MAC protocols are implemented with a strict layer separation so that their operation is transparent to the user and very lightly coupled with the Physical layer

through a physical convergence layer. In Proximity-1, because of the need for manual intervention during fault recovery, the MAC entity is tightly coupled with both the user as well as the Physical layer.

The Proximity-1 MAC sublayer has two functional components:

a) connection control – activation and deactivation of a Proximity-1 session and token pass for half-duplex operations;

b) configuration control – setting configuration such as frequency channel, data rate, modulation, encoding, and duplex mode.

The MAC sublayer has an extensive set of connection and configuration control functionalities that are not traditionally associated with the MAC sublayer concept. Many of these functionalities are directly offered to the Proximity-1 user as a service. Please refer to 3.2.2 ('Control Plane Services') for more detail.

3.3.5.2   Connection Control

The most important connection control function of the MAC sublayer is the hailing procedure by which two spacecraft can establish both Data Link and Physical layer connections. Another connection control function is the token passing process used in half-duplex mode. The MAC sublayer performs the necessary control function of regulating data flow and toggling the transceivers between transmit and receive.

3.3.5.3   Configuration Control (Comm Change)

Configuration control allows Proximity-1 to change both Physical layer parameters such as frequency, coding, data rate, modulation, as well as Data Link layer parameters such as Data Link layer protocol and duplex mode, from the default configuration established through hailing. Changing duplex mode, however, requires the initiator to change to the new duplex state before the directive is transmitted across the Proximity link. In summary, reconfiguration is a key component in operating efficiently and optimizing operations in a heterogeneous environment where there is a wide variety of spacecraft with different communication capabilities.

3.3.5.4   State Variables

The MAC sublayer maintains and updates several state variables that relate to the status and configuration of a Proximity-1 session. The MODE variable controls whether the Proximity-1 connection is *inactive*, within the transmitting or listening phase of the hailing procedure (*connecting-T*, *connecting-L*), or ready to provide data services (*active*). The TRANSMIT variable controls whether the transmitter is turned on and ready to send out Proximity-1 frames. The DUPLEX parameter selects full, half, simplex-transmit, or simplex-receive operations.

3.3.5.5   Persistence Activity

A persistent activity is a process for ensuring reliable communication between a caller and a responder using the Expedited QOS while transmitting exclusively from the MAC queue. Because of the potential for frame loss due to corruption across the space link, these MAC control activities require a persistence process to ensure that supervisory protocol directives are received and acted upon correctly. Persistence activities may be linked in series to accomplish a task, but persistence applies to only a single activity at a time. Proximity-1 (reference [8]) currently defines three persistent activities: Hailing, i.e., Session Establishment (see reference [8], subsection 6.2.4 and tables 6-8 and 6-11); COMM_CHANGE (see reference [8], subsection 6.2.4 and tables 6-9 and 6-12); and Resynchronization (see reference [8], subsections 7.1.3.2 and 7.1.3.3).

Operationally, when the persistence signal is set to *true*, all other data flow is suspended except for this persistent activity operated out of the MAC sublayer until one of the following conditions occurs: 1) the expected RESPONSE to the activity is received within the allocated response time, i.e., WAITING_PERIOD; 2) the expected RESPONSE is not received within the WAITING_PERIOD, but before the LIFETIME of the activity has expired, so that the persistent action can be repeated; 3) the RESPONSE is not received within the LIFETIME and the activity is terminated without success. In any case, the protocol notifies the vehicle controller about the result of the activity and resets the Persistence signal to *false*.

3.3.6   PHYSICAL LAYER

3.3.6.1   General

The primary functions of the Physical layer are:

a) Establish a communications channel upon which the data can flow. This process includes configuration of the following Physical layer parameters: frequency, data rate, modulation, and coding parameters, such that common operating characteristics exist in both spacecraft.

b) Acquire the carrier signal.

c) Once the carrier is acquired, acquire bit-level synchronization using the acquisition sequence (see reference [10], subsection 3.3.1). Note that the idle sequence was so chosen as to enable acquisition in the presence of data on the Proximity link.

d) Provide the tail sequence for a configurable fixed period prior to terminating transmission in order for the receiving spacecraft to process the received data unit fully (required for convolutional decoding and bit lock assurances).

e) Maintain a synchronous data link when data is unavailable for transmission by inserting the idle sequence onto the physical channel.

f)   Provide **Carrier_Acquired** and **Bit_Inlock_Status** signals to the MAC sublayer.

Note that the acquisition, idle, and tail sequences utilize the same bit pattern for all of these functions, i.e., 352EF853 (in hexadecimal).

3.3.6.2   Bit-Level Synchronization

Proximity-1 handles bit-level synchronization in two different ways. When message exchanges are brief and there is no intention of establishing long-term communication on a channel, the Physical layer operates in a message-switching fashion, acquiring carrier and bit-level synchronization for each message and relinquishing synchronization when there is no message to send. Since synchronization is not maintained between transmissions of messages, Proximity-1 defines this as an 'asynchronous channel'. If message transmissions are sporadic and have long signal-period dead times, this approach can provide significant energy savings.

However, if there is a significant amount of data to be transmitted, Proximity-1 will establish a physical channel that operates in a circuit-switched fashion: once carrier lock and bit synchronization are achieved, bit synchronization will be maintained by inserting an idle sequence when necessary. Because bit-level synchronization is maintained, Proximity-1 defines this as a 'synchronous channel'. This has the advantage of reducing acquisition overhead, but may still result in higher energy consumption if there are many idle periods during the session.

3.3.6.3   Full- and Half-Duplex Operations

Proximity-1 supports full-duplex communication, where physical channels are established for both the forward and the return links. This approach will minimize overhead associated with channel acquisition since it is nominally only done once during link establishment. However, for half duplex, the physical channel must be re-established every time the link is turned around, a fact that may contribute to significant overhead, if there are many turnarounds during a session.

In Proximity-1, hailing in both full- and half-duplex operations as well as half-duplex sender-receiver turnarounds are regulated by four parameters: **Carrier_Only_Duration, Acquisition_Idle_Duration**, **Tail_Idle_Duration**, and **Hail_Wait_Duration**. They are assigned values in the MIB to optimize the operation of the link establishment process. To overcome the uncertainty in Doppler shift and timing, the caller sends a carrier-only signal to establish carrier lock with the responder; this signal is then followed by an acquisition idle sequence to achieve bit-level synchronization. After the hail directives are transmitted, the idle sequence functions as a tail sequence sent to ensure that all of the data will be processed through the receiver chain and in particular through the convolutional decoder, if applicable. The **Hail_Wait_Duration** provides a period of time for the sender's receiver to listen for any response on the return channel without interference from the transmitter.

3.3.6.4 Frequency Allocation/Assignment

The Proximity-1 Physical layer has provision for Proximity operations at the following frequency bands: UHF, S, X, and Ka. At the time of this document's publication, only frequency channel assignments at UHF band have been made. Assignments at the other frequency bands remain TBD until a user need is identified for them.

The UHF frequency allocation consists of 60 MHz between 390 MHz to 450 MHz. The forward frequency band is defined from 435 to 450 MHz. The return band is defined from 390 to 405 MHz. There is a 30 MHz dead band between them. There are 16 channel assignments; 8 of these channels so far have been assigned a specific frequency or frequency range.

Proximity-1 has the concept of two channel types: hailing and working channels. The hailing channel is used only for the initial link establishment. Time spent on the hailing channel should be kept to a minimum to allow other users access to it. It is recommended that after link establishment through hailing is accomplished, one transition to the working channel (if available) as soon as possible and continue to conduct operations on it.

The working channel is used for the bulk of Proximity communication which includes any reacquisitions (due to loss of lock or half-duplex turnarounds), nominal command, telemetry, and timing data services and link termination.

The frequencies assigned for hailing are enterprise specific. This enables an enterprise to pick the forward and return frequencies for hailing to best meet the needs of the individual missions within that enterprise. Note that if the Proximity link radio equipment supports only a single channel (i.e., a single forward and return frequency pair), then the hailing channel will be the same as the working channel. This is the case for both the NASA Mars Odyssey Orbiter and the ESA Mars Express Orbiter for which channel 0 represents both the hailing and working channel. However, if interoperability among enterprises is required at UHF, the Recommended Standard has reserved channel 1 (435.6 MHz Forward, 404.4 MHz Return) as the hailing channel. Note that this frequency pair was chosen to minimize the amount of bandwidth occupied by the hailing channel for two reasons: first, because hailing is done at a low data rate and therefore is a low bandwidth activity and, second, because the use of the remaining spectrum is maximized for working channels.

For more complex configurations with one or multiple forward frequencies paired with one or multiple return frequencies, the Recommended Standard allows for up to eight additional forward frequencies and eight additional return frequencies to be selected as long as they are distinct from all previous assignments.

3.3.6.5 Bit Error Rate

For asynchronous data link operations, the Bit Error Rate (BER) on the link should be less than or equal to 1×10^{-6}. The Physical layer also needs to meet specific Doppler tracking and acquisition requirements (see reference [10] subsection 3.5.6).

4   PROTOCOL OPERATIONS

4.1   MESSAGING BETWEEN TWO TRANSCEIVERS

4.1.1   GENERAL

There are several types of messages that can flow between two transceivers during a Proximity-1 session. Figure 4-1 shows the flow through the communication protocol stack of various Proximity-1 data and control messages from one Proximity-1 node to another.
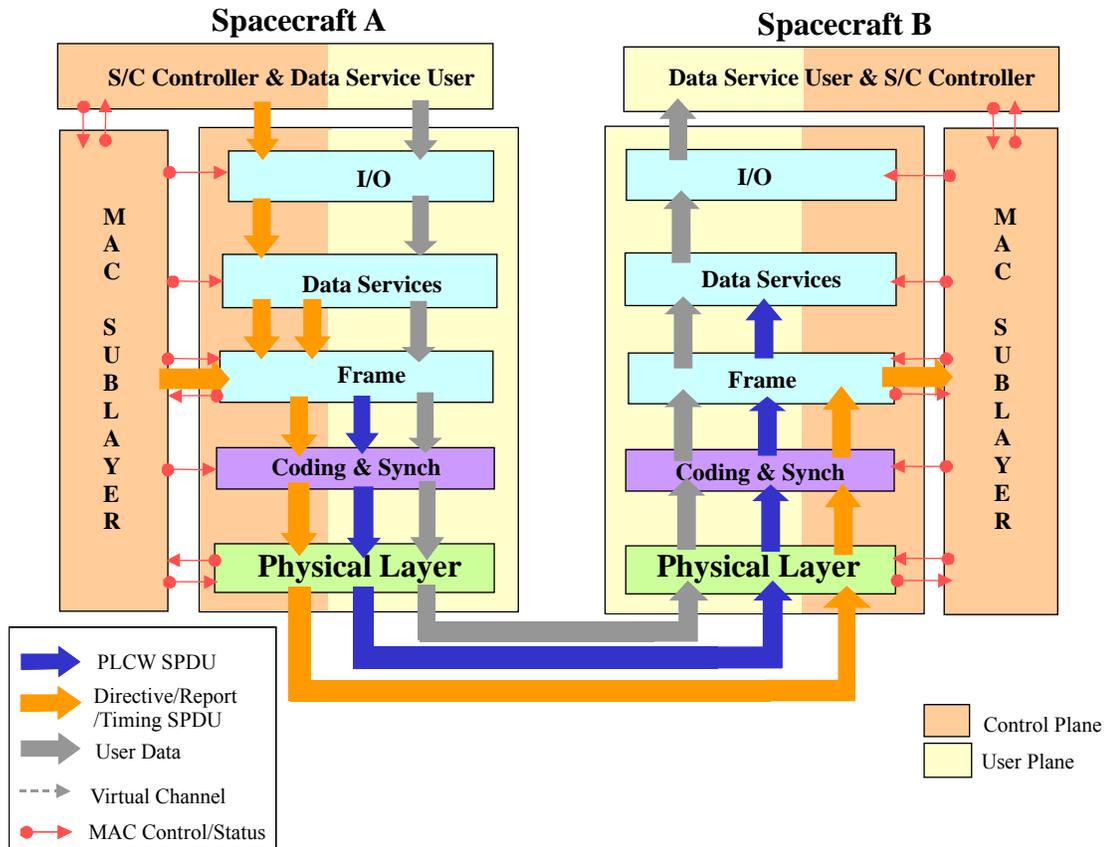
Figure 4-1: Flow of Proximity-1 Data and Messages

The user data always originates from the sender-user and passes through all the sublayers except the MAC sublayer. The previous section showed how user data is submitted as an SDU, then converted to frames and PLTUs, and then modulated onto the physical channel as a bit stream. The receiver reverses the process to recover the original data SDUs.

For control messaging such as directives, either the user or the MAC sublayer may issue an SPDU, but it will all be delivered to the remote entity's MAC sublayer, which will then decode the directives and perform the requested actions. Another control message exchanged between two Proximity-1 nodes is the Proximity Link Control Word (PLCW). These are used for acknowledging frame receipt and making retransmission requests. PLCW flow from the data receiver's Frame sublayer back to the data sender's Data Services sublayer.

4.1.2 HAILING PROCEDURE

Hailing is an interactive process by which two Proximity-1 entities can discover each other and establish an asynchronous data link over a physical connection between them. Hailing can be accomplished by means of a demand or by negotiation. In a demand hail, a fixed set of directives is transmitted in a Proximity-1 frame in order to configure the partnered transceiver. A negotiated hail allows for a more flexible link establishment process based upon requests/reports for capabilities and status. In either case, the initiator of this process is referred to as the caller, and the intended communication partner of the caller is referred to as the responder. The hailing procedure utilizes the process of a 'persistence activity' to ensure reliable handshaking between the caller and the responder. The key idea behind a persistence activity is that it is transactional: an action is declared complete only when the expected response is observed within a reasonable delay.

Before the hailing process begins, the caller will load a P-frame that contains the hailing message (one or more Proximity-1 directives) into the MAC sublayer's transmission queue. The hail frame will be addressed to the intended responder and contains directives such as the SET TRANSMITTER PARAMETERS and SET RECEIVER PARAMETERS directives for configuring the responder's transmitter and receiver. Then by setting the MODE parameter to *connecting-T*, the hail frame will be radiated on the hailing channel. The responder's transceiver is configured by its vehicle controller to be in the *connecting-L* mode. Once it correctly receives the hail frame or acquires carrier lock and bit synchronization,[6] it responds back with a status report frame and begins data service. Upon receipt of this frame, the initiator resets persistence to *false* and begins data service. See figure 4-2, Proximity-1 Link Establishment.

---

[6] Implementation choice per the Proximity-1 Recommended Standard.
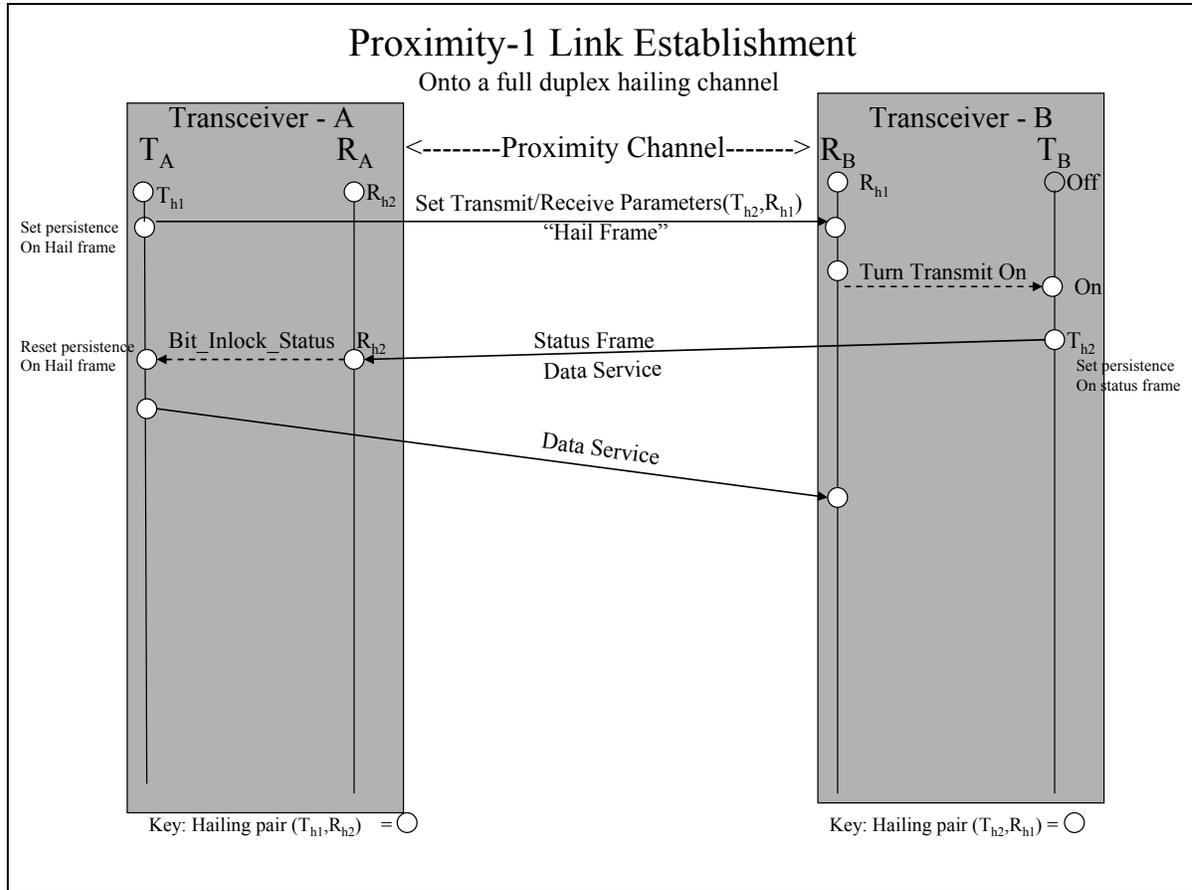
Figure 4-2: Proximity-1 Link Establishment

4.1.3 FULL DUPLEX: MOVING OFF THE HAILING CHANNEL AND ONTO A WORKING CHANNEL

Figure 4-3 shows the steps involved in moving onto a Working Channel immediately after the link is established on the Hailing Channel. In this case, the hail frame also contains parameters to change the responder's transmit and receive frequency. Upon successful receipt of the hail, the responder changes receive frequency to R_{w1} and turns on its transmitter. It transmits a status frame back to the initiator on its hail frequency. Upon receipt of the hail response, the initiator resets persistence on the hail activity and sets its receiver to the new receive frequency, R_{w2} and new transmit frequency, T_{w1}. The initiator now in data services transmits frames. Upon acquiring bit synchronization or a valid frame, the responder changes its transmitter to T_{w2} and begins data service on the working channel.

Proximity-1 Link Establishment
Onto a working full duplex channel via a demand hail
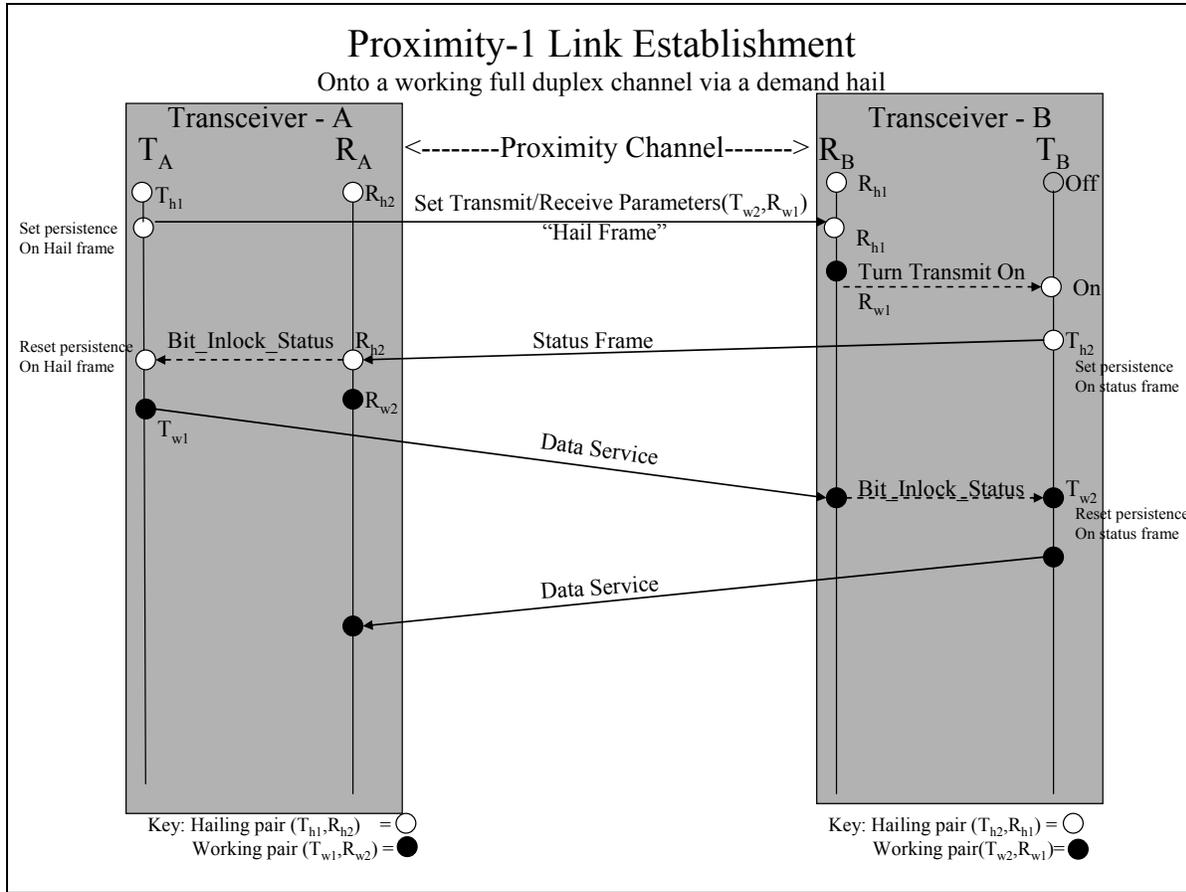
Figure 4-3: Moving onto the Working Channel in Full Duplex

4.1.4   LINK ESTABLISHMENT ONTO A HALF-DUPLEX CHANNEL

For the half-duplex option, whether the session eventually uses a working channel or stays on the hailing channel, a *three-way handshake* is necessary to ensure the integrity of the hailing process. The hail frame and subsequent response frames both contain the SET CONTROL PARAMETERS directive that has the Token indicator. This indicator basically gives the other entity control of the channel. If such an indication is not given at the end of each entity's transmission, then either the loss of bit-lock or a half-duplex timer overrun will have the same effect of reversing the transmit-receive orientation.[7] See figure 4-4, Half-Duplex Link Establishment, below.

---

[7] See reference [8], subsection 4.3.10.6.3 for details.

Figure 4-4: Half-Duplex Link Establishment

4.1.5   COMMUNICATION CHANGE DURING DATA SERVICES

During a session (after link establishment) it is possible to change Physical layer parameters such as data rate, modulation, frequency, duplex mode, etc. Typically the decision to change communications configuration is made by the caller, but either node can initiate the change. If the prototypical orbiter-to-lander scenario is considered, the orbiter, which is the caller, can be seen to be the natural candidate for making configuration decisions; as the receiver of data, it will be able to assess the channel condition on the return channel (from surface to orbit) based on received SNR and other metrics derived from Proximity-1 operations.

Figure 4-5: Generic Full-Duplex Communications Change

Figure 4-5 depicts a possible sequence of events for a typical communications change scenario using a full-duplex channel.

One can also reconfigure communication parameters during data services on a half-duplex channel.

4.1.6 DATA SERVICES PHASE

4.1.6.1 General

The main feature of the data service phase is the option to multiplex two grades of data service onto the link: Sequence Controlled and/or Expedited. The Expedited services attempts delivery of an SDU only once without guarantee of correct reception. The Sequence Controlled service uses a go-back-*n* ARQ to guarantee correct reception. The following subsections describe the internal operation of Proximity-1 with more emphasis on the Sequence Controlled service in more detail.

4.1.6.2   Flow of Proximity-1 Frames On the Sending Side

4.1.6.2.1   I/O Sublayer: Frame Construction

Once a session enters the data services phase, the user may supply SDUs that carries either user data or protocol directives. These are then converted to Proximity-1 transfer frames called U-frame and P-frame, respectively. The header of each frame will contain the sequence number, the QOS, physical channel, Port ID, and source/destination Spacecraft ID, PDU type, and other information necessary to identify and provide the appropriate grade of service to the SDUs contained in these frames. Separate sequence numbers are kept for each QOS.

Once a frame is generated, it is separately queued based on the QOS requested. The Sequence Controlled queue holds frames that require the Sequence Controlled service; the Expedited queue holds frames that require the Expedited service. Note that frames are queued strictly in the order in which they are generated.

4.1.6.2.2   Data Services Sublayer: ARQ Error Control

Since the Data Services sublayer tracks the state of the ARQ algorithm, it maintains two queues to handle possible re-transmissions: a 'New Sequence Controlled Frame Queue' and a 'Sent Sequence Controlled Frame Queue', which stores Sequence Controlled frames that are sent until they are positively acknowledged. Depending on the state of the ARQ algorithm, one frame from either queue will be loaded into a 'Sequence Controlled Frame Buffer' for delivery. For Expedited service the Data Services sublayer maintains no queue, but just a buffer that holds a single Expedited frame extracted directly from the Expedited queue in the I/O sublayer.

4.1.6.2.3   Frame Sublayer: Prioritized Multiplexing

The Frame sublayer performs prioritized multiplexing of frames for transmission across the physical channel. Transfer frames can come from several sources, and they are given priority as expressed in 3.3.3.

After the Frame sublayer selects a frame to be sent, the Coding and Synchronization sublayer pre-pends an ASM, computes a CRC, appends it to the frame and delivers this PLTU to the Physical layer. There, the PLTU is optionally convolutionally encoded and transmitted over the physical channel.

4.1.6.3   Flow of Proximity-1 Frames on the Receiving Side

4.1.6.3.1   Gerenal

On the receiving end of a Proximity-1 session, the Physical layer will decode the optionally convolutionally encoded bits from the physical channel and submit them to the Coding and Synchronization sublayer as a bit stream.

4.1.6.3.2   Frame Delimiting and Transmission Error Detection

The decoded bit stream contains frames separated by synchronization markers and possibly idle bits inserted by the Physical layer. The Frame sublayer will look for the ASM, extract the frame, and decode the CRC parity check. If there is no error, then the frame is accepted. Otherwise it is rejected and the NEED_PLCW parameter is set to true.

4.1.6.3.3   Validate and Route Frames

When the Frame sublayer receives frames from the Coding and Synchronization sublayer, it performs another frame validation process based on the format and content of the frame. It checks that the correct frame format is used (Version 3) and the destination Spacecraft ID is valid for the current Proximity-1 session. (Optionally tests whether the source Spacecraft ID is valid when Test_Source is *true*).

For frames that are accepted, SPDUs are extracted from P-frames. If an SPDU is a PLCW, it is forwarded to the Data Services sublayer so that the ARQ states can be proper updated; the other types of SPDUs are sent to the MAC sublayer. The U-frames are forward to the Data Services sublayer.

4.1.6.3.4   Frame Sequencing

The Data Services sublayer, upon receipt of a frame, will determine whether there is a missing frame and whether retransmission needs to be requested. For Sequence Controlled services, a missing frame is detected when there is a gap between the sequence numbers between the current frame and the last received frame on the same physical channel; a request for retransmission can be done by setting the NEED_PLCW parameter to *true*. For Expedited service, the received frame is passed directly on to the I/O sublayer.

4.1.6.3.5   SDU Extraction

The Data Services sublayer will deliver U-frames to the I/O sublayer in order, but possibly with gaps, when using Expedited service. The I/O sublayer will handle the task of extracting the SDUs contained in the frames, reassembling them when needed, and delivering them to the user through the designated port. The I/O sublayer must deliver only complete SDUs to the user; therefore missing frames that contains piece(s) of one or more SDUs will cause them to be discarded.

4.1.6.4   COP-P: A Go-Back-*N* ARQ Algorithm

4.1.6.4.1   General

The Sequence Controlled service is an instance of a go-back-*n* ARQ implementation, which in Proximity-1 is called COP-P because of its similarity to the implementation in

reference [3], Communications Operation Procedure-1 (COP-1). As depicted in figure 4-6, COP-P describes the interaction between a sender-receiver pair. The procedure on the sending side is called FOP-P, and its counterpart of the receiving side is called FARM-P.

Unbalanced Point-to-Point Configuration
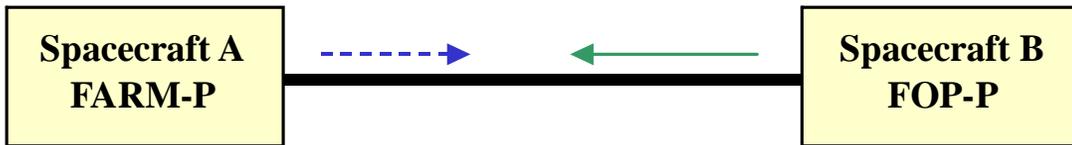
Balanced Point-to-Point Configuration
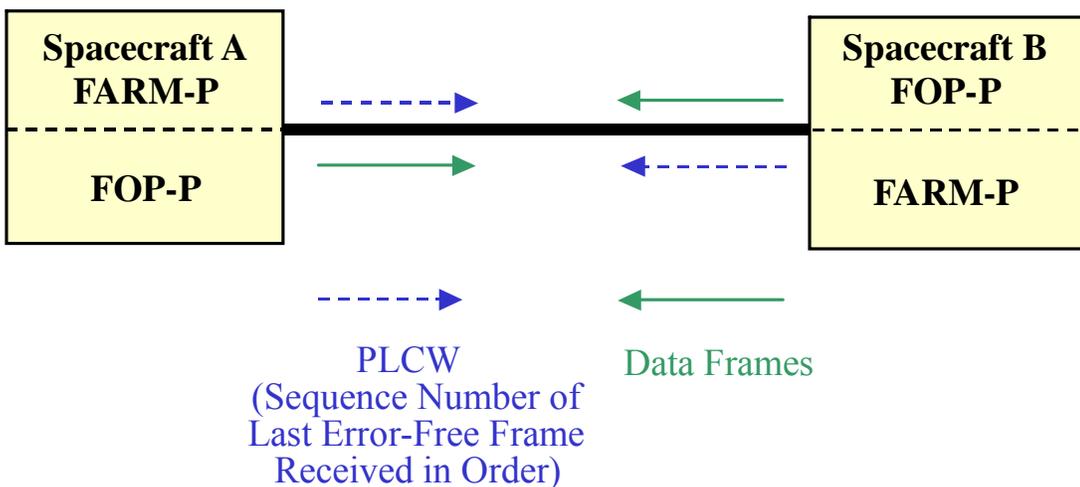
Figure 4-6: COP-P Configurations

Functionally, FOP-P controls the order of transmission of frames, and this algorithm resides in the Data Services sublayer. The FARM-P process is implemented by the functionalities on the receiving side of the Coding and Synchronization, Frame, and Data Services sublayers. It checks each received Sequence Controlled frame for order of arrival and frame duplication. Only unique and in-order frames are accepted. The FARM-P process sends feedback messages called PLCWs to the FOP-P process regarding the sequence number of the last frame received in order. Using this information, the FOP-P process can decide whether retransmission is necessary and which frames can be cleared from its buffer space. The interaction between a FOP-P/FARM-P pair implements a go-back-n ARQ algorithm.

4.1.6.4.2   PLCW Generation

Because COP-P stores all state information such as the sequence number of the most recently sent and acknowledged frame in the Data Services sublayer, one naturally assumes the Data Services sublayer would also create the actual feedback messages, or PLCW. However, Proximity-1 directly generates each PLCW in the Frame sublayer using counter values provided by the FARM-P process. This has several benefits:

  a) eliminates queuing delays by generating each PLCW right before it is sent out for coding and transmission on the physical channel,

  b) assures each PLCW carries the most up-to-date information when it is generated, thereby improving half-duplex operation by allowing just one PLCW to be generated for each turnaround.

4.1.6.4.3   COP-P State Diagram

The following are two abridged state diagrams for the FOP-P and FARM-P processes. Not all counter and state variables are included, and some details are omitted in order to provide a more descriptive presentation of the COP-P state tables. The state diagrams below show the state transitions for Sequence Controlled service on a single physical channel.

Event 4: Sync Timer Expired
Action: If Resync Local is true
1. Initiate Set V(R) Activity
2. Transition to S2

Event 0: Initialize
Action:
1. Initialize State Variables
2. Start Sync Timer

Event 1: Frame sublayer needs frame to transmit
Action:
1. Test to see that the appropriate frame is sent to the frame sublayer.
2. Adjust VV(S) and or V(S) accordingly

Event 2: Received a valid PLCW while in S1
Action:
1. Clear buffer of acknowledged frame.
2. Clear Sync Timer.
3. If ReTx requested, set VV(R)=N(R)
4. Store PLCW

Event 5: Set Transmission_Window Request (accept, set)

Event 6: Set Synch_Timeout Request (accept, set)

Event 7: Reset Request (accept, initialize)

Event 8: Invalid Request (reject)

Event 3: Received an Invalid PLCW
Action:
1. Set VV(S)=NN(R)
2. Start Sync_Timer

Figure 4-7 : FOP-P State Diagram

Event 0: Initialize
Action:
    1. Set all counters to zero
    2. $R(S) = $ *false*
    3. Set $NEED_PLCW = $ *true*

Event 6: Valid Seq. Ctrl frame
          $N(S) < V(R)$
Action: Discard Frame

Event 1: Invalid Frame Arrives
Action: Discard frame

Initialize   0

Event 3: Valid Expedited Frame Arrives
Action: Accept Frame/Increment Expedited
          Frame Counter

Event 4: Valid Seq. Ctrl Frame Arrives
           In Sequence $N(S) = V(R)$
Action:
    1. No Retransmission needed, so set
        $R(S)=$*false*
    2. Increment counter value $V(R)$
    3. Trigger feedback message by
        setting $NEED_PLCW = $ *true*

Event 7: Received request for
           PLCW content
Action: Report $R(S), V(R)$, and
          Expedited Frame Counter
          (to the Frame sublayer)

ACTIVE

1   3   6   7   4   2   5

Event 2: Received a SET
           $V(R)$ directive
Action:
    1. Set $V(R)$ to the specified
        value
    2. Set $R(S)$ to *false*
    3. Trigger feedback
        message by setting
        $NEED\_PLCW=$*true*

Event 5: Valid Seq. Ctrl Frame Arrives
           Gap detected $N(S) > V(R)$
Action:
    1. ReTx is needed, so set $R(S)=$*true*
    2. Trigger feedback message by
   setting $NEED_PLCW = $
    3. Reject the frame.

Figure 4-8 : FARM-P State Diagram

4.1.6.5   Time Sequence Diagram

4.1.6.5.1   General

The FOP-P and FARM-P processes interact through the exchange of Sequence Controlled U-frames, directives, and PLCWs, which can acknowledge one or multiple accepted frames. The time sequence diagram uses a shorthand notation to show the value and content of each message. A PLCW is designated by a pair of brackets with three fields: **[PC,V(R),R(S)]**; directives are denoted by curly brackets: **{directive}**; Sequence Controlled U-frames are designated by parenthesis: **(PC,N(S))**.

Note that $V(R)$ is the sequence number (plus one) of the last Sequence Controlled frame accepted by the receiver; $R(S)$ is a binary flag used to request retransmission. PC is the Physical Channel ID (either '0' or '1'), and $N(S)$ is the sequence number of a Sequence Controlled frame. Also note that the value of the **SYNCH_TIMER** is constantly adjusted in a fashion similar to a 'sliding window', whose size is the **Synch_Timeout** parameter.

Proximity-1 allows each physical channel to have its own COP-P process, however; concurrent COP-P procedures are required and the reporting is then required to contain the status for each physical channel. The specification does not address data prioritization and multiplexing into the output bit stream of simultaneous multiple physical channel operations.

The following time sequence diagrams show operations over one physical channel using up to two COP-P processes. For simplicity of illustration only, a go-back-two implementation is shown.

4.1.6.5.2 Full-Duplex Operation

This subsection shows examples of message exchanges under full-duplex operation:

– Unbalanced configuration (see figure 4-6):

a) One COP-P Process – data frame error (loss) recovery (figure 4-7)'

b) One COP-P Process – PLCW error (loss) recovery (figure 4-8).

Spacecraft A   **Spacecraft B**

(0,0)

(0,1)

[0,1,No]   Frame 0 accepted

(0,2)

[0,2,No] (0,3)   Frame 1 accepted

(Frame 3 has error; it is discarded)

[0,3,No]   Frame 2 accepted

Frame 4 is valid but discarded

(0,4)   due to *out-of-sequence delivery*

[0,3,Yes   Frame 3 is negatively ac

Frame 3 is resent

(0,3)

(0,4)   Frame 4 is sent simultaneously to fill up the window size (N=2)

[0,5,No]   Frames 3 and 4 are accepted
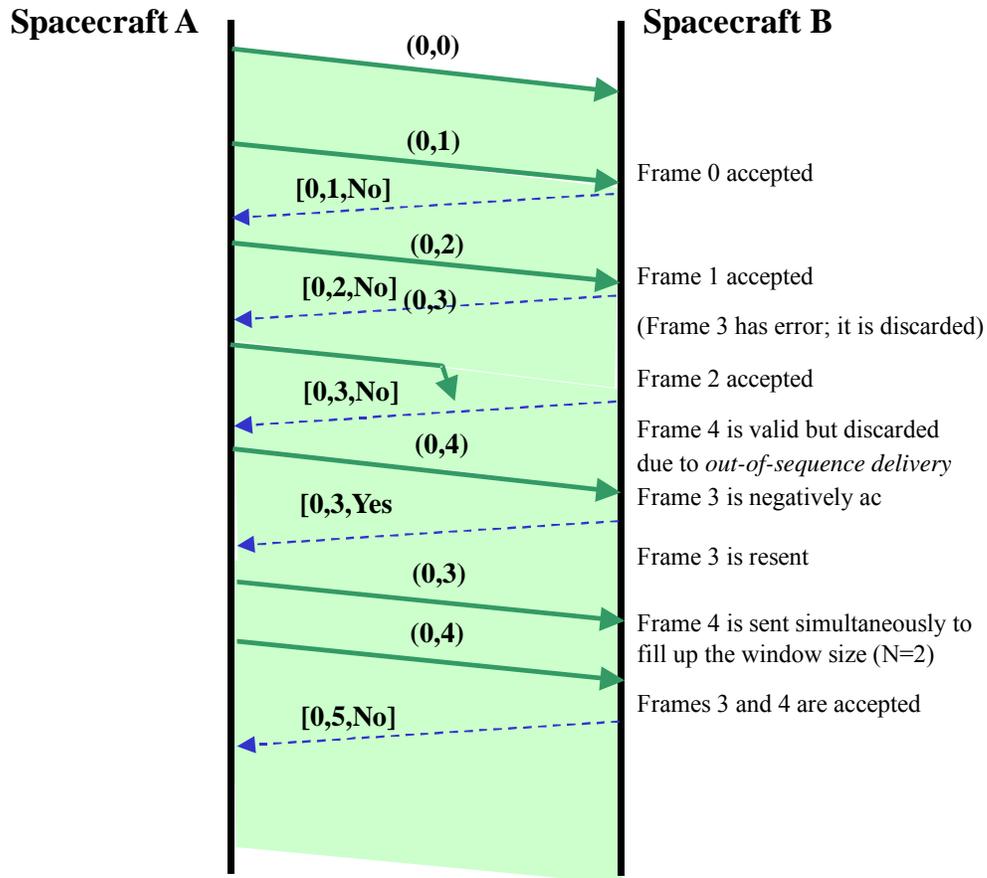
Figure 4-9: Recovery from Error Frame (Go-Back-Two)

Figure 4-7 show the typical error recovery process implemented in Proximity-1. Note that when a data frame is lost or dropped due to a transmission error, no PLCW is generated because there is ambiguity about the identity of the rejected frame. If a PLCW were generated to report this event, the FOP-P process would not be able to decide with confidence whether retransmission is necessary because it has to guess the identity of the rejected frame. Therefore, the FARM-P process will only generate a PLCW when a valid data frame arrives *out-of-sequence*, in which case the receiver, the sender, and which frames are missing are clear.
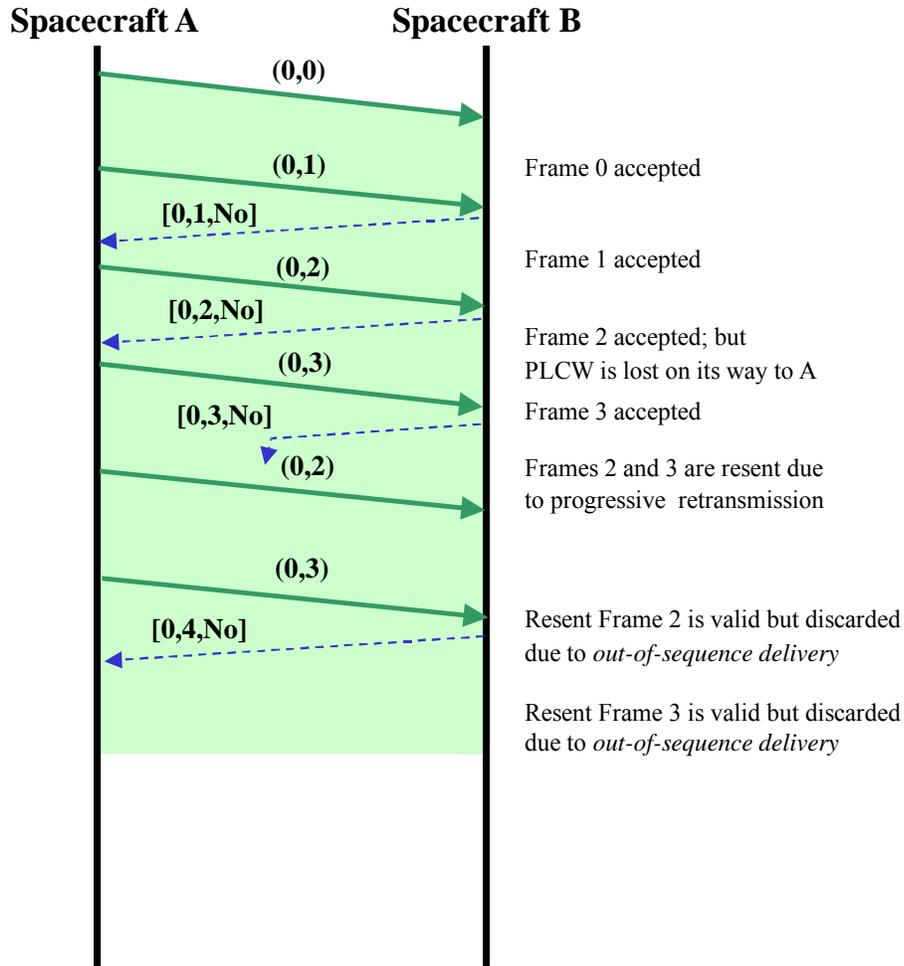
Figure 4-10: Recovery from Lost PLCW

Figure 4-8 shows the recovery process when a PLCW is lost or rejected due to a transmission error. The time sequence diagram shows that losing a PLCW is no more complex than losing a data frame because the sender is waiting for acknowledgement and re-transmitting frames that the receiver believes have already been acknowledged. This scenario works because the sender continues to retransmit the frames it has in its Sent queue until it receives a PLCW or until the SYNCH_TIMER expires.

4.1.6.5.3   Half-Duplex Operation

The most critical aspect of half-duplex operation is the choice of window size, timeout value, and turnaround mechanism. While for full-duplex operation, the optimal window size is one that just allows the sender to fill the pipe continuously, in half-duplex mode the channel error rate and turnaround overhead must also be added to the equation. One can expect that under good channel conditions or long turnaround time (include round-trip time and hardware considerations), the transmission window size should be large. This strategy is consistent with the prototypical unbalanced link, where the majority of the data flow is from the surface asset to the relay orbiter. Another difference of half-duplex operations is the Synch_Timeout parameter: because acknowledgement cannot occur until the receiving node is allowed to transmit on the channel, it must be adjusted accordingly.

In Proximity-1, the turnaround can be triggered by two methods: (1) sending SET CONTROL PARAMETERS directive with the Token indicator, and (2) automatically switching roles when the half-duplex timer overflows. Although using the timer make the process fairly automatic and the use of a directive unnecessary, one needs to use a long guard band duration to avoid cutting off transmission and/or reception in the middle of a frame. *Therefore, it is suggested that the explicit use of the directive should be the primary turnaround mechanism, while the timer is used only as a backup in case the directive is not received.*
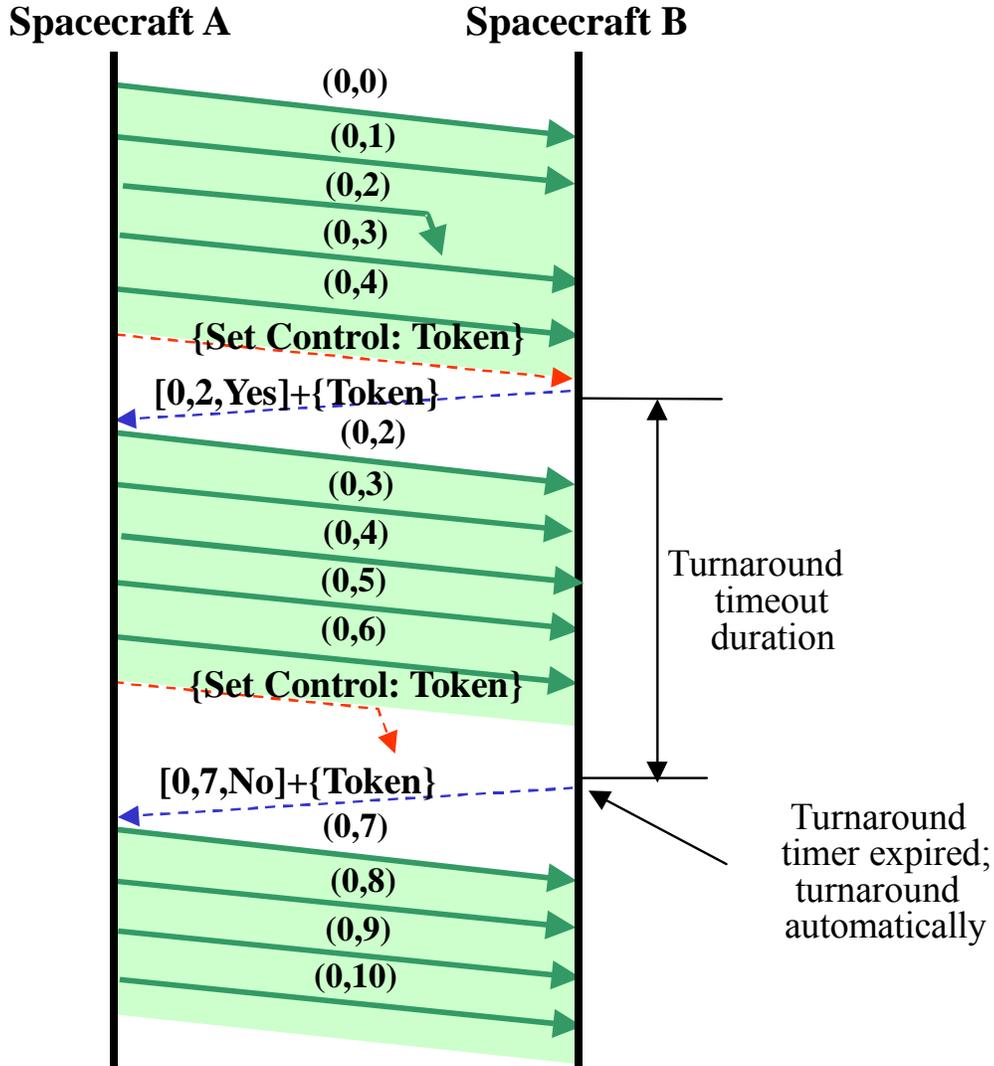
Figure 4-11: Half Duplex Running a Single COP-P Instance

The time sequence diagram in figure 4-11 illustrates how only one acknowledgement is necessary for each turnaround as well as how the turnaround timer is used as a safety backup in case the turnaround directive is not received.

4.1.6.6   Session Termination

4.1.6.6.1   Use of the LNMD and RNMD Directives

A graceful termination requires that some kind of handshaking at the end of the session occurs. In Proximity-1 (reference [8]), *a two-way handshake* is described. Figure 4-12 shows the general termination process between two spacecraft.

The first step is that the lander's vehicle controller has provided the last SDU that it wishes to send to the Proximity-1 I/O sublayer. At this point the local vehicle controller provides a local LNMD directive to the local transceiver of Spacecraft A. As a result of receiving this command, Transceiver A generates a RNMD (remote no more data to send) directive and transmits it across the Proximity link. In effect, this directive is telling the partnered transceiver that Spacecraft A has no more data to send. Similarly, when Vehicle Controller B provides its last SDU to the Proximity-1 I/O sublayer, it generates a local LNMD directive and sends it to Transceiver B. Thereafter, Transceiver B generates a RNMD directive informing the partnered transceiver that it has no more data to send and shortly thereafter begins the link termination process including terminating the carrier signal and turning its transceiver off.
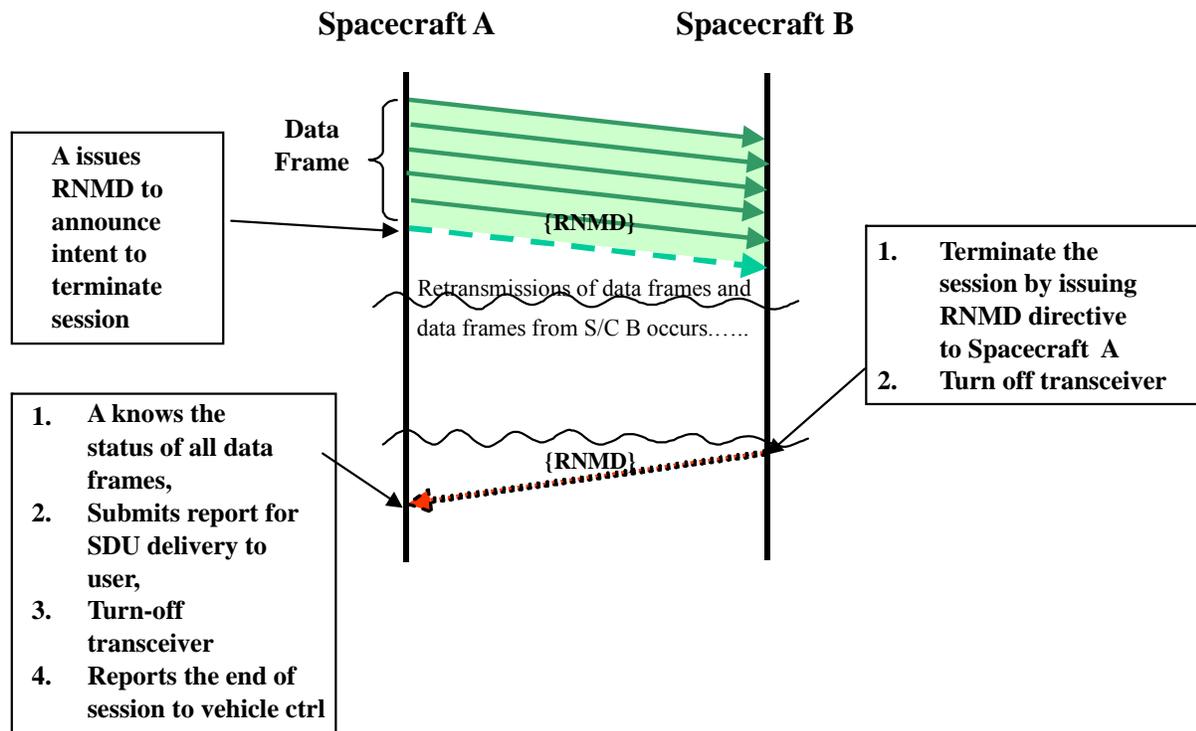
Figure 4-12: Session Termination (Two-Way Handshake)

4.1.6.6.2   Carrier-Loss

Of course there is the possibility either that a graceful termination is not needed, or even not possible (e.g., the spacecraft already went out of view or is malfunctioning), or that the RNMD directive is lost or corrupted. As a fail-safe measure, Proximity-1 will terminate a session when the carrier loss timer expires. It is up to the local vehicle controller to decide whether it is necessary to establish a new session.

4.1.6.6.3 Continuity of Data Transmission

It is possible that when a session has ended, data transmission is not yet complete from the user's perspective. In such a case, a graceful termination is important so that the sender has the necessary information to know which frames, and therefore which SDUs, have been received *completely and correctly*. This allows for a seamless resumption of SDU transfer at the beginning of the next session. Because Proximity-1 delivers only complete SDUs to the user, any SDU that is not received completely will not be delivered to the user; therefore the entire SDU must be sent again in the next session.
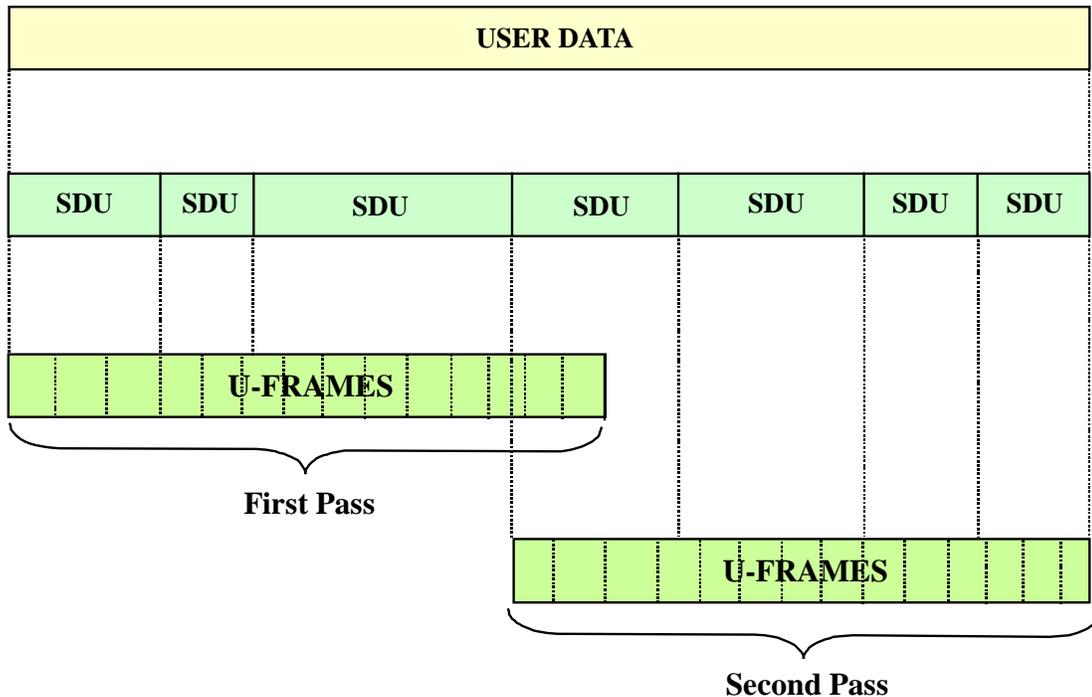
Figure 4-13: Transmission of User Data Across Two Sessions

It is possible for the follow-on session to resume after the last acknowledged frame. To do so requires the vehicle controller to do some extra accounting of frame numbers associated with the SDU transferred between the two sessions. There will be a discontinuity in the sequence count between the end of the first pass and the start of the second pass. The vehicle controller needs to reconcile this difference in order to verify that the SDU is complete. However, such an operation is more complicated than the standard procedure for providing complete SDUs to the user in the protocol.

4.2 OPERATIONAL ISSUES

4.2.1 GENERAL

In the following subsections a few issues are raised only as suggestions for users or implementers of Proximity-1.

4.2.2 LARGE SDU SPANNING MORE THAN 256 FRAMES

Proximity-1 tracks the delivery status of each frame by its frame sequence number, which is assigned by the I/O sublayer and stored in an eight-bit field in the Proximity-1 header. When the user supplies an SDU, it can be transferred over the link by sending the contents over one or multiple frames if necessary. Therefore, Proximity-1 can track the delivery progress of each SDU by keeping a record of the last frame sequence number associated with it. Complications may arise when an SDU is so large that it occupies more than $2^8=256$ frames, within which duplicated frame sequence numbers exist. A more sophisticated tracking algorithm could be used to track sequence count rollovers; an alternative approach would be to segment the SDU. Proximity-1 segmentation keeps track of which segment is the first, middle, and last segment associated with an SDU, so that the SDU can be reassembled from those segments on the receive side.

4.2.3 WINDOW SIZE AND RETRANSMISSION TIMEOUT VALUES

The performance of the go-back-*n* ARQ algorithm depends strongly on the choices of the window size. The overall goal is throughput efficiency but different methods are applied to derive the optimal setting depending upon the channel type.

For a full-duplex channel in which data is transmitted only in one direction, throughput efficiency is independent of how often acknowledgements are sent because there is a dedicated channel for that purpose. Therefore, the optimal window size is one that just allows the sender to fill the physical channel with data frames. This can be derived from the round-trip propagation time, transmission time of a data frame, and the delay for the processing, generation, and transmission of the acknowledgement. The case when data is transmitted in both directions is slightly more complicated, because both data and acknowledgements (i.e., PLCWs) must share the link. In this case, Proximity-1 provides a set of frame prioritization rules that guarantee that the link will be equally shared between them.

For half-duplex operations, the process is more complicated because of turnarounds. The turnaround frequency is determined by the operations concept for data transfer. Fundamentally, missions with unbalanced links, e.g., high-rate telemetry links and low-rate command links, will minimize the number of turnarounds. The use of a higher-layer selective repeat protocol like CFDP over Expedited Proximity-1 is recommended.

4.2.4 BIDIRECTIONAL VS. UNIDIRECTIONAL DATA TRANSFER EFFECT ON LINK THROUGHPUT

4.2.4.1 Bidirectional User Data Transfer

Since Proximity-1 supports simultaneous bidirectional data transfer, U-frames can flow simultaneously in both the forward and return link directions. When using the Sequence Controlled QOS in both directions, these links will contain at a minimum both P-frames (PLCWs) and U-frames. For the purpose of this discussion, the return link is considered to be the primary link and the forward link to be secondary (the mission operations concept allows link assignments to be reversed). The amount of user data traffic flowing in one direction, e.g., the forward link, will affect the throughput of user data flowing in the opposite direction. In effect, for bidirectional data transfer, transferring user data in one direction reduces the bandwidth available for the PLCW traffic supporting the acknowledgements in the opposite direction. Without receiving PLCWs in a timely fashion, the return link, once it reaches its maximum transmission window size, N, is prohibited from sending any new Sequence Controlled U-frames. At this point, the return link becomes less than 100-percent efficient, since the transceiver will progressively retransmit the frames in its Sent queue until it receives a positive acknowledgement enabling it to move its transmission window forward.

4.2.4.2 Unidirectional User Data Transfer

Proximity-1 supports unidirectional user data transfer, i.e., the transfer of U-frames exclusively on either the forward or the return link. Unidirectional data transfers are inherently more efficient in terms of maximizing data throughput in one direction. However, data transferred in the other direction incurs a latency penalty resulting in the combination of forward and return user data transfer's taking longer to complete. In particular, when using the Sequence Controlled QOS, the node receiving the U-frames transmits only P-frames but no U-frames; i.e., the return traffic is restricted to PLCWs (per COP-P) and the optional transmission of MAC frames used to transmit directives across the link.

4.2.4.3 User Controlled Proximity-1 Parameters Affecting Data Throughput Efficiency

The Proximity-1 user-set parameters (both on the forward and return data links) that affect the user data throughput efficiency of both forward and return Proximity-1 links are:

  a) data rates (using asymmetrical forward and return links);

  b) frame lengths;

  c) maximum transmission window size, i.e., N.

Other factors that affect Proximity-1 link data throughput are:

  a) channel frame error rates;

  b) light time signal propagation delays;

  c) delay components in the transmit and receive chains of the transceiver.

On a bidirectional link, 100-percent throughput efficiency degrades when the fastest link has sent all of its data (within the window size) while the slow link completes sending a U-frame before it sends the acknowledging PLCW (one PLCW can acknowledge a window length of frames). During this time the fastest link starts progressively retransmitting the data from the Sent queue, since the window size has been reached. In order to maintain 100-percent throughput on the fastest link and minimize the non-progression of this link, the user frame size on the slowest link needs to be reduced (see figure 4-14).
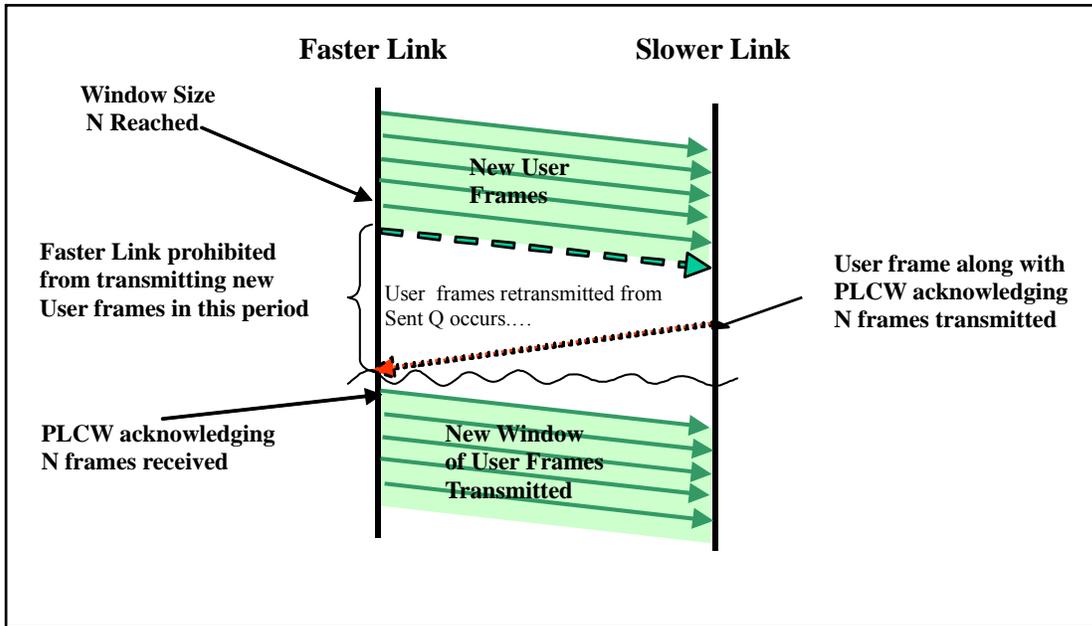
Figure 4-14:Loss of 100-Percent Throughput on the Faster Link of a Bidirectional Link

Figure 4-15 shows that by reducing the frame size on the link with the slowest data rate, 100-percent data throughput is maintained on both links even as the fastest link achieves the highest Proximity-1 data rate. This assumes a window size (N) of 127, which is the maximum allowed in Proximity-1.

Largest Frame Size Allowed to Maintain 100% Throughput, N=127
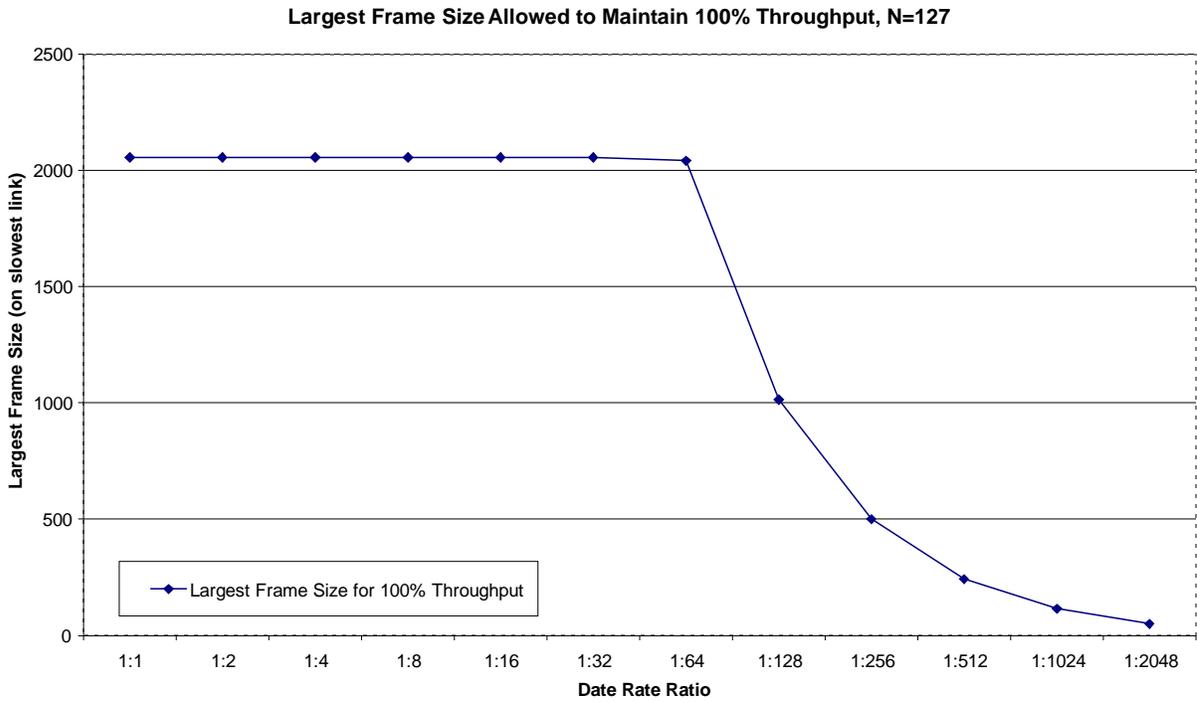
Figure 4-15:Largest Frame Size (Octets) on Slowest Link Allowed to Maintain 100-Percent Throughput (*N*=127)

Largest Frame Size Allowed to Maintain 100% Throughput, N=1
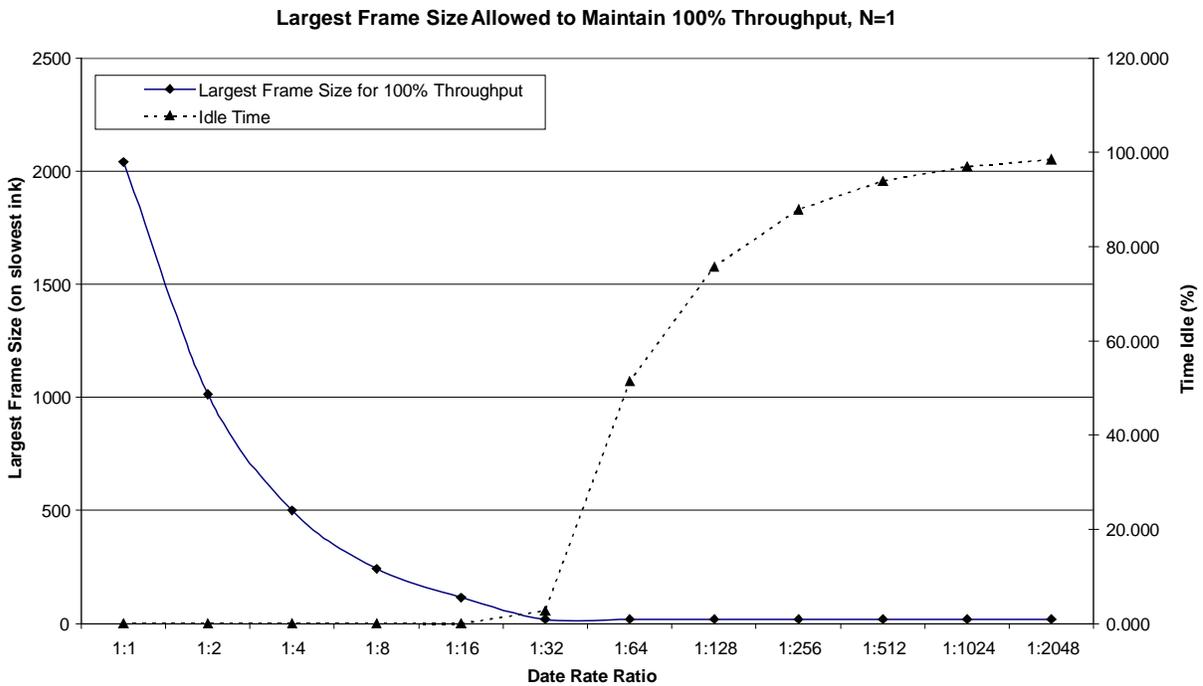
Figure 4-16:Largest Frame Size (Octets) on Slowest Link Allowed to Maintain 100-Percent Throughput (*N*=1)

In figure 4-16, showing the extreme case of N=1, it can be seen that the ability to maintain 100-percent data throughput at a data rate ratio of 1:32 is lost. This results in the faster side of the link's spending more time progressively retransmitting frames. This loss of throughput is proportional to the data rate ratio, slower link to faster link rate. The higher this ratio, the longer the faster link fails to make progress in transmitting new user frames. For completeness, figure 4-17 shows N=32 and the loss of data throughput of 100 percent at a data rate ratio of 1:512. The data rate ratio at which we lose data throughput of 100 percent is proportional to the transmission window size. So the bigger the window size, the larger the data rate ratio can be before losing data throughput of 100 percent.
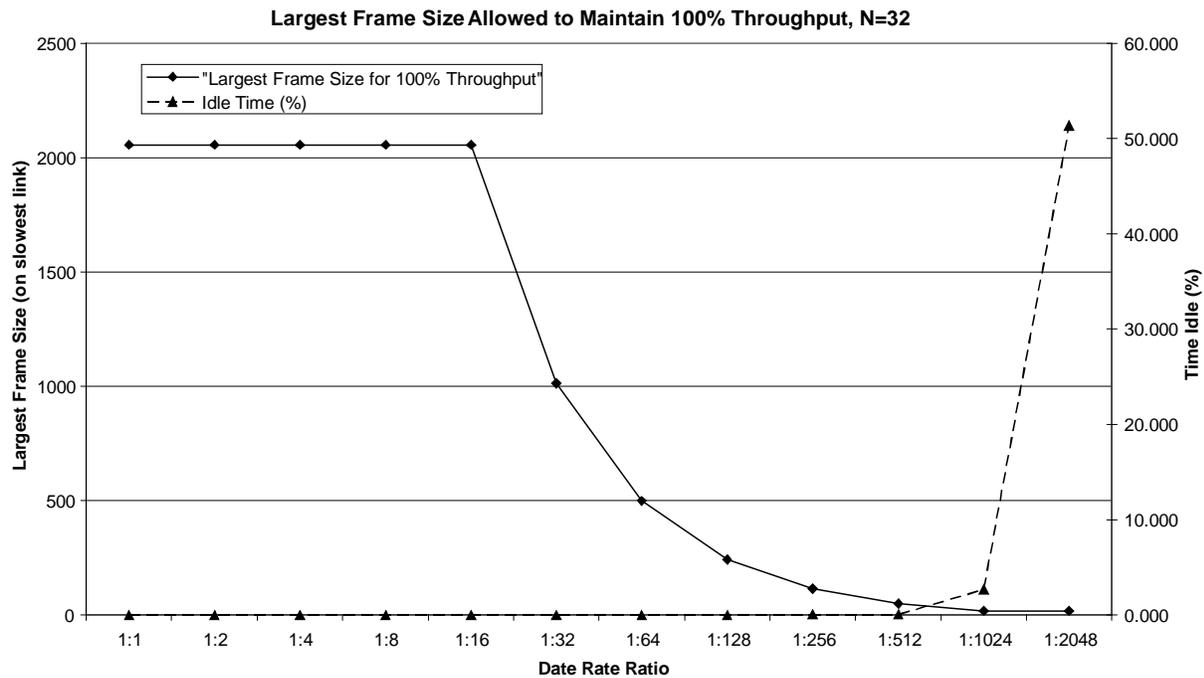
Figure 4-17: Largest Frame Size (Octets) on the Slowest Link Allowed to Maintain 100-Percent Throughput (N=32)

4.2.5 MULTIPLEXING SERVICE DATA UNITS OVER THE PROXIMITY-1 LINK

Proximity-1 can support the multiplexing of multiple protocol SDUs into the data link of a physical channel. For example, it is possible to multiplex CCSDS Space Packets, IPv4 or IPv6 datagrams, and SCPS datagrams using the same or distinct Port IDs of a given physical channel. The packet version number contained within each SDU is used to identify the packet and route it to the appropriate upper-layer protocol.

4.2.6   UNRELIABLE BIT STREAM MODE

Since the Proximity-1 protocol architecture allows the user to couple directly with the Physical layer, the vehicle controller can directly supply a bit stream to the Physical layer for transmission, thus bypassing Proximity-1 framing. Although this is not part of the Recommended Standard, it is easily implemented within the Proximity-1 architectural framework.

5   SCENARIOS

5.1   ASSUMPTIONS

This section uses the prototypical surface asset-to-orbiter scenario to illustrate how Proximity-1 operates under a wide variety of topologies. This specific scenario is chosen because it is the most relevant scenario for near-term Proximity-1 operations. To understand scenarios other than surface-to-orbiter communications and those that involve balanced configuration where data flows in both directions, the reader can generalize the surface asset-to-orbiter scenario by considering the orbiter as the caller, and/or receiver of data; similarly one can generalize the lander as the sender of data, and/or the responder in the hailing process. Because Proximity-1 has the flexibility such that the master is not necessarily the caller and the receiver of data, the number of possible operational configurations is large. The reader can make the necessary extensions.

5.2   STATIC VS. DYNAMIC ENVIRONMENT

Proximity-1 is designed to operate across a spectrum of scenarios and configurations. At one end of the spectrum is a static environment where Proximity-1 will operate between a fixed pair of spacecraft that usually belongs to the same mission. The caller-responder roles between these spacecraft are usually predefined and, even when change is allowed, are pre-scheduled. Configuration and operational parameters are limited and pre-selected to match the capability of the spacecraft. Also the frequency band is predefined and dedicated. Usually in such cases, the hailing channel is also used for data transmission. A static environment allows implementers to study a given scenario exhaustively and fine-tune Proximity-1 to operate under a fixed but optimized configuration.

At the other end of the spectrum is a dynamic environment where multiple spacecraft from different missions may communicate with each other to exchange data, query information, or form cooperative teams to achieve scientific objectives dynamically driven by unforeseen events and new mission objectives added after launch. The operations concept can be highly dynamic because the number of potential communication partners is high, and their corresponding configurations are less predictable. It is no longer cost effective to dedicate a frequency channel to each possible pairing; rather, radio spectrum must be shared to ensure interoperability while remaining adaptive. In such a case, there is a commonly designated hailing channel used for all possible spacecraft pairs; data transmission will typically take place on a separate pair of frequency channels.

5.3   TOPOLOGY

5.3.1   GENERAL

Proximity-1 is basically a single master topology, like HDLC and other Data Link layer protocols, because its frame header can only address either the source or the destination node, but not both. In order to operate in a multicast or many-to-many topology where each node is treated equally, both the source and destination addresses must be included.

Under a single master topology there are two possible topological variations: (1) one-to-one and (2) one-to-many.

5.3.2 ONE TO ONE

5.3.2.1 General

A one-to-one topology includes all scenarios where each spacecraft is only in communication with one other spacecraft. There can be multiple one-to-one Proximity-1 links operating at the same time. In a static environment, each one-to-one Proximity-1 link operates completely independently of others; in a dynamic environment, there will be some interaction, especially during the link establishment process.

5.3.2.2 One-to-One Link using a Dedicated Channel

Step 1. Orbiter sends hail to lander on predefined dedicated channel.

Step 2. While orbiter maintains carrier signal, the lander responds to the original hail. After bit lock or valid frame is received, connection is established.

Step 3. Lander starts data transmission; configuration is unbalanced.

Step 4. Lander notifies orbiter there is no more data to send.

Step 5. Orbiter notifies lander there is no more data to send; orbiter turns off transceiver; lander receives directive and turns off transceiver as well.

Figure 5-1: One-to-One Link Using a Dedicated Channel

When multiple one-to-one links exist, they each operate independently because each link has its own dedicated frequency channel pair.

5.3.2.3 Single One-to-One Link using Separate Hailing and Working Channels

Typically, in a dynamic environment, Proximity-1 uses separate hailing and working channels. Sharing a common hailing channel allows any spacecraft to make contact with any other spacecraft in that enterprise; switching to a separate working channel is necessary because this releases the shared hailing channel for use by other spacecraft.

Step 1. Orbiter sends hail to lander on shared hailing channel; lander instructed to use working channel 1.

Step 2. Lander responds to the original hail on return hailing channel and orbiter listens on working channel.

Step 3. If no interference, orbiter switches to working forward channel; after bit lock or valid frame is received at the orbiter, the link is established.

Step 4. Lander begins data transmission on return working channel.

Step 5. Lander notifies orbiter there is no more data to send.

Step 6. Orbiter notifies lander there is no more data to send; orbiter then turns off transceiver; lander receives directive and turns off transceiver.

Figure 5-2: One-to-One Link Using Separate Hailing and Working Channels

5.3.2.4 Multiple One-to-One Links in a Dynamic Environment: Contention on the Hailing Channel

When the hailing channel is shared, then contention resolution is important. This subsection describes how contention is resolved by using random back-off and a carrier sensing mechanism.

First it is assumed there are two orbiters and two landers. Orbiter 1 wants to establish a one-to-one link with lander 1; orbiter 2, with lander 2. However, both orbiters decide to initiate the hailing process at almost the same time so that the carrier sensing mechanism does not detect the other's presence. Two scenarios are shown: (1) both hailing messages collided and (2) only one hailing message is corrupted.

The next case examined is one in which one hailing process starts before the other one so that a carrier sensing mechanism allows the second caller to back-off its hailing attempt and avoid collision.
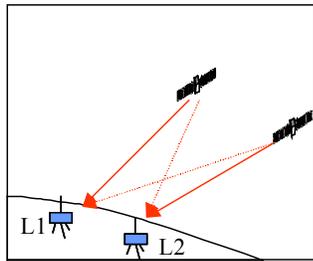
Step 1. Orbiter 1 and orbiter 2 send hail at about the same time; both hailing messages are corrupted due to collision.
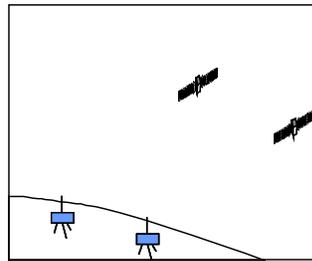
Step 2. No lander responds; orbiters back-off for random period of time before next attempt.
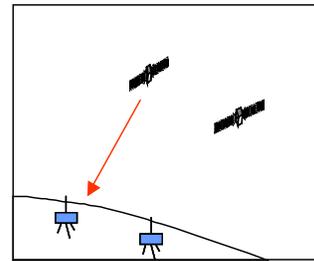
Step 3. Orbiter 1 timer expires so it starts hailing first; orbiter 2 still in back-off mode…Hail is successful.
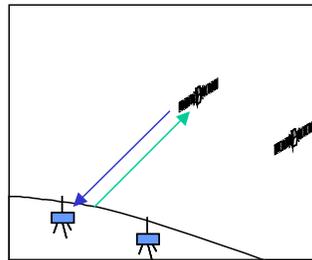
Step 4. Orbiter 2 backed off while Orbiter 1 established a prox-1 connection on a separate working channel.
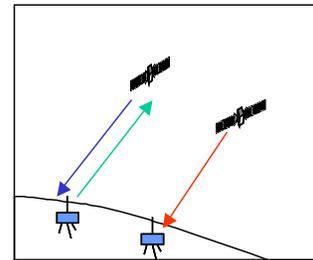
Step 5. Orbiter 2 hails successfully.

Figure 5-3: Simultaneous Hailing; Both Landers in View of the Orbiters
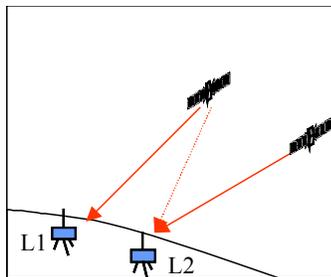
Step 1. Orbiter 1 and orbiter 2 send hail at about the same time; but L1 is not in footprint of orbiter 2, so hail is successful for L1.
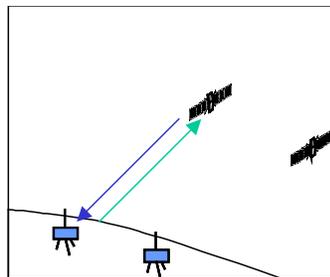
Step 2. Orbiter 2 backed off while Orbiter 1 established a prox-1 connection on a separate working channel.
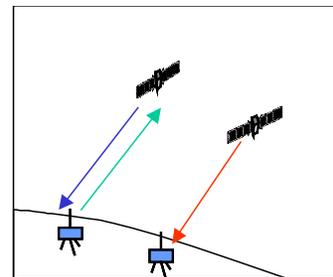
Step 3. Orbiter 2 hails successfully without collision.

Figure 5-4: Simultaneous Hailing; Only One Lander in View of Both Orbiters

Step 1. Orbiter 1 and lander 1 are exchanging hailing and response messages; orbiter 2 senses carrier signal on one or both forward and return hailing channels. So it delays its own hailing process.

Step 2. Orbiter 2 backed off while Orbiter 1 established a prox-1 connection on a separate working channel.

Step 3. Orbiter 2, after backing off for random period of time, hails successfully without collision.

Figure 5-5: Collision Avoidance Using Carrier Sensing on Hailing Channel

5.3.2.5 Multiple One-to-One Links in a Dynamic Environment: Contention on the Working Channel

In all the scenarios described above, when contention on the hailing channel is resolved, the result is that one lander-to-orbiter link will be established before the second caller can start its hailing process. At this point, the second caller needs to perform additional carrier sensing in order to determine which working channel is available for use. Such a decision must be made before hailing because it is part of the hailing message.

Step 1. Orbiter and lander1 have established prox-1 link on channel 2; orbiter sweeps channel and hears carrier signal on either one or both the forward or return link on channel 2.

Step 2. Orbiter 2 selects channel 3 as working channel to avoid collision with the existing prox-1 link; it issues hailing message instructing lander 2 to use channel 3.

**Step 3**. Orbiter 2 established prox-1 link on Channel 3.

Figure 5-6: Collision Avoidance on the Working Channel

5.3.3   ONE-TO-MANY EXTENSIONS

For the one-to-many topology, a single orbiter can communicate with multiple landers at the same time using Proximity-1.

On the forward (command) link, an orbiter can transmit commands to multiple landers simultaneously, because a lander can discriminate the commands intended for itself based upon the Spacecraft ID field in the Proximity-1 transfer frame header. In addition, broadcasting of commands to all or a group of landers within the footprint of an orbiter is also possible using the Spacecraft ID as a global or group identifier.

For the return link, if the relay orbiter has only a single channel receiver available, then the orbiter can time share the link. Since the process of link establishment and moving to a working frequency takes time, the orbiter overflight can be envisioned as divided into a subset of smaller sessions each dedicated to a specific lander as opposed to a time division multiplex scheme. If the relay orbiter has multiple transceiver resources available, frequency assignments may be created that allow simultaneous accommodation of multiple surface assets. One possible configuration is to use multiple surface-to-orbit and multiple orbit-to-surface frequency pairs in FDMA/FDM fashion, running full-duplex Proximity-1 protocols on each frequency pair independently. A higher-layer process would be needed to manage these links, particularly if the orbiter uses its same SCID for all of them.

A more efficient approach is to consider multiple return links using FDMA, and a single forward channel (used to multiplex PLCW responses to the surface). All landers would need to tune to the same receive channel. Again, several full-duplex Sequence Controlled Proximity-1 sessions could be supported as long as adequate margin is available in the shared forward link.

ANNEX A

ACRONYMS

ARQ	Automatic Repeat Queuing
ASM	Attached Synchronization Marker
BER	Bit Error Rate
BNSC	British National Space Centre
CCSDS	Consultative Committee for Space Data Systems
CFDP	CCSDS File Delivery Protocol
COP-P	Communication Operations Procedure for Proximity links
DFC ID	Data Field Construction Identifier
ESA	European Space Agency
FARM-P	Frame Acceptance and Rejection Mechanism Proximity
FOP-P	Frame Operations Procedure-Proximity
IPv4	Internet Protocol Version 4
IPv6	Internet Protocol Version 6
MAC	Medium Access Control
MER	NASA Mars Exploration Rover
MIB	Management Information Base
MRO	NASA Mars Reconnaissance Orbiter
NASA	National Aeronautics and Space Administration
N(R)	Last acknowledged frame sequence number +1
PCM	Pulse Code Modulation
PDU	Protocol Data Unit
P-Frame	Supervisory/Protocol Frame
PLCW	Proximity Link Control Word

PLTU	Proximity Link Transmission Unit
PSK	Phase Shift Keyed
QOS	Quality of Service
RF	Radio Frequency
R-S	Reed-Solomon
SAP	Service Access Point
SCID	Spacecraft Identifier
SCPS-NP	Space Communications Protocol Standards-Network Protocol
SDU	Service Data Unit
U-frame	User Data Format
UHF	Ultra High Frequency
V(S)	Value of Sequence Number