

Consultative Committee for Space Data Systems

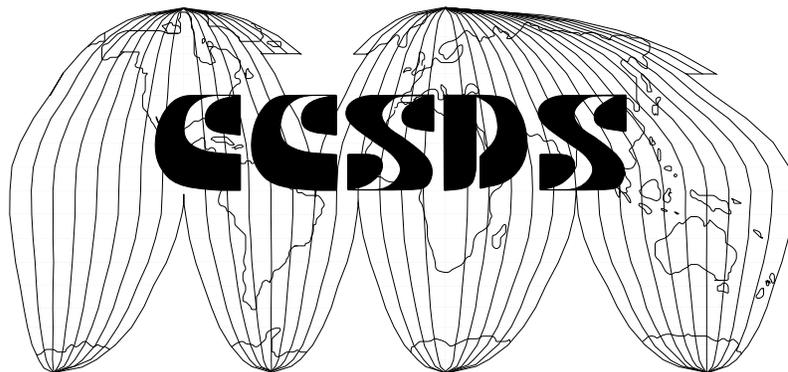
**DRAFT REPORT CONCERNING SPACE
DATA SYSTEM STANDARDS**

**SPACE COMMUNICATIONS
PROTOCOL SPECIFICATION (SCPS)—
RATIONALE, REQUIREMENTS,
AND APPLICATION NOTES**

CCSDS 710.0-G-0.3

DRAFT GREEN BOOK

April 1997



AUTHORITY

Issue:	Draft Green Book, Issue 3
Date:	April 1997
Location:	N/A

This document has been approved for publication by the Management Council of the Consultative Committee for Space Data Systems (CCSDS) and represents the consensus technical agreement of the participating CCSDS Member Agencies. The procedure for review and authorization of CCSDS Recommendations is detailed in reference [1], and the record of Agency participation in the authorization of this document can be obtained from the CCSDS Secretariat at the address below.

This document is published and maintained by:

CCSDS Secretariat
Program Integration Division (Code MG)
National Aeronautics and Space Administration
Washington, DC 20546, USA

FOREWORD

This document is a technical Report for use in developing space communication networks and has been prepared by the Consultative Committee for Space Data Systems (CCSDS). The space communication protocols described herein are intended for use within missions that are cross-supported between Agencies of the CCSDS.

Through the process of normal evolution, it is expected that expansion, deletion or modification to this document may occur. This Report is therefore subject to CCSDS document management and change control procedures which are defined in Reference [1].

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

Member Agencies

- British National Space Centre (BNSC)/United Kingdom.
- Canadian Space Agency (CSA)/Canada.
- Centre National d'Etudes Spatiales (CNES)/France.
- Deutsche Forschungsanstalt für Luft- und Raumfahrt e.V. (DLR)/Germany.
- European Space Agency (ESA)/Europe.
- Instituto Nacional de Pesquisas Espaciais (INPE)/Brazil.
- National Aeronautics and Space Administration (NASA)/USA.
- National Space Development Agency of Japan (NASDA)/Japan.
- Russian Space Agency (RSA)/Russian Federation.

Observer Agencies

- Australian Space Office (ASO)/Australia.
- Austrian Space Agency (ASA)/Austria.
- Central Research Institute of Machine Building (TsNIIMash)/Russian Federation.
- Centro Tecnico Aeroespacial (CTA)/Brazil.
- Chinese Academy of Space Technology (CAST)/China.
- Communications Research Laboratory (CRL)/Japan.
- Danish Space Research Institute (DSRI)/Denmark.
- European Organization for the Exploitation of Meteorological Satellites (EUMETSAT)/Europe.
- European Telecommunications Satellite Organization (EUTELSAT)/Europe.
- Federal Service of Scientific, Technical & Cultural Affairs (FSST&CA)/Belgium.
- Hellenic National Space Committee (HNSC)/Greece.
- Indian Space Research Organization (ISRO)/India.
- Industry Canada/Communications Research Centre (CRC)/Canada.
- Institute of Space and Astronautical Science (ISAS)/Japan.
- Institute of Space Research (IKI)/Russian Federation.
- KFKI Research Institute for Particle & Nuclear Physics (KFKI)/Hungary.
- MIKOMTEK: CSIR (CSIR)/Republic of South Africa.
- Korea Aerospace Research Institute (KARI)/Korea.
- Ministry of Communications (MOC)/Israel.
- National Oceanic & Atmospheric Administration (NOAA)/USA.
- National Space Program Office (NPSO)/Taiwan.
- Swedish Space Corporation (SSC)/Sweden.
- United States Geological Survey (USGS)/USA.

DOCUMENT CONTROL

Document	Title	Date	Status
CCSDS 710.0-G-0.2	Report Concerning Space Data Systems Standards: SCPS Concept and Rationale	August 1996	Draft Green Book— Second Issue
CCSDS 710.0-G-0.3	Report Concerning Space Data Systems Standards: SCPS Concept and Rationale	April 1997	Draft Green Book— Third Issue

CONTENTS

<u>Section</u>	<u>Page</u>
1 INTRODUCTION	1-1
1.1 PURPOSE AND SCOPE.....	1-1
1.2 DOCUMENT STRUCTURE.....	1-2
1.3 DEFINITIONS	1-3
1.4 REFERENCES	1-4
2 OVERVIEW	2-1
2.1 BACKGROUND	2-1
2.2 RATIONALE FOR SCPS.....	2-3
2.3 APPLICABILITY OF SCPS	2-4
2.4 SCPS LAYERS/PROTOCOLS OVERVIEW	2-4
2.5 SCPS TESTING AND VALIDATION.....	2-5
2.6 SUMMARY	2-6
3 RATIONALE FOR SCPS	3-1
3.1 CHANGING ENVIRONMENT	3-1
3.2 OPERATIONAL CONSTRAINTS	3-1
3.3 EXISTING PROTOCOLS	3-2
3.4 SUMMARY	3-5
4 SCPS NETWORK PROTOCOL	4-1
4.1 SCPS-NP OVERVIEW	4-1
4.2 TERMINOLOGY	4-2
4.3 SCPS-NP REQUIREMENTS	4-3
4.4 SHORTCOMINGS OF USING IP IN SPACE NETWORKS	4-5
4.5 CAPABILITIES OF THE SCPS NETWORK PROTOCOL.....	4-6
4.6 CONFIGURATION ALTERNATIVES	4-8

CONTENTS (CONTINUED)

<u>Section</u>	<u>Page</u>
5 SCPS SECURITY PROTOCOL	5-1
5.1 SCPS-SP OVERVIEW.....	5-1
5.2 SCPS-SP HERITAGE.....	5-3
5.3 MISSION APPLICATIONS OF SCPS-SP	5-3
5.4 SCPS-SP PROTECTION METHODS.....	5-4
5.5 SECURITY ASSOCIATION ATTRIBUTES.....	5-5
5.6 SECURITY PROTOCOL OPERATION	5-5
5.7 END-SYSTEM TO INTERMEDIATE-SYSTEM INTERACTIONS	5-6
6 SCPS TRANSPORT PROTOCOL (SCPS-TP)	6-1
6.1 SCPS-TP OVERVIEW	6-1
6.2 SCPS-TP REQUIREMENTS	6-2
6.3 SCPS-TP SERVICES.....	6-6
6.4 SCPS-TP EXTENSIONS TO TCP.....	6-6
6.5 SPACE MISSION APPLICATIONS OF SCPS-TP.....	6-12
6.6 SCPS-TP SUMMARY	6-13
7 SCPS FILE PROTOCOL (SCPS-FP)	7-1
7.1 OVERVIEW OF SCPS-FP.....	7-1
7.2 SCPS-FP SERVICES	7-2
7.3 SCPS-FP REQUIREMENTS	7-3
7.4 SCPS-FP MODIFICATIONS TO FTP	7-3
7.5 CONFORMING IMPLEMENTATIONS OF THE SCPS-FP	7-5
ANNEX A GLOSSARY	A-1
ANNEX B ACRONYMS.....	B-1
ANNEX C FREQUENTLY ASKED QUESTIONS	C-1
ANNEX D PROTOCOL FUNCTIONAL REQUIREMENTS [PFR].....	D-1
ANNEX E SCPS BENT-PIPE EXPERIMENT REPORT (SUMMARY)	E-1
ANNEX F SCPS FP/STRV 1B FLIGHT EXPERIMENT (SUMMARY)	F-1

CONTENTS (CONTINUED)

<u>Figures</u>	<u>Page</u>
1-1 SCPS Layered Protocols	1-1
1-2 Layered Protocol Model.....	1-3
2-1 Extension of Internet into Space.....	2-1
2-2 SCPS end-to-end Services	2-2
4-1 SCPS Network Protocol (SCPS-NP)	4-1
4-2 Encapsulating Network Configuration.....	4-8
4-3 Translating Network Configuration.....	4-9
5-1 SCPS Security Layer	5-1
5-2 Security Header Layout.....	5-4
5-3 SCPS-SP Operation.....	5-5
5-4 SCPS-SP in an On-board Intermediate System.....	5-7
6-1 The SCPS Transport Protocol.....	6-1
7-1 SCPS File Transfer Protocol	7-1
7-2 Internet FTP Model.....	7-2
<u>Tables</u>	<u>Page</u>
4-1 Support of SCPS Network Requirements by IP	4-5
4-2 Support of SCPS Network Requirements by SCPS-NP	4-6
5-1 Security Services.....	5-2
5-2 Space Mission Applications of SCPS-SP.....	5-3
6-1 Factors Affecting TCP Performance in Non-Terrestrial Environments	6-3
6-2 SCPS TP Transport Services	6-6
6-3 SCPS-TP Modifications to TCP to Address Communication Problems	6-7
6-4 Space Mission Applications of SCPS-TP	6-12

1 INTRODUCTION

1.1 PURPOSE AND SCOPE

This Report describes the rationale, use, and intent of a set of layered space/ground protocols, the "Space Communications Protocol Standards" (SCPS). As shown in Figure 1-1, the SCPS protocol layers are specified in a set of four CCSDS Recommendations (Reference [3-6]). The SCPS protocols support the transfer of space mission data through space-to-ground and space-to-space data subnetworks. These protocols are not intended for transfer of space mission data that occurs wholly within ground systems, but rather are focused on the unique requirements of data transfer through subnetworks that involve a space data transmission path. The SCPS can be used as an integrated protocol stack, or the individual protocols can be used in combination with CCSDS or internet protocols to create custom profiles to support the requirements of particular Agencies or missions.

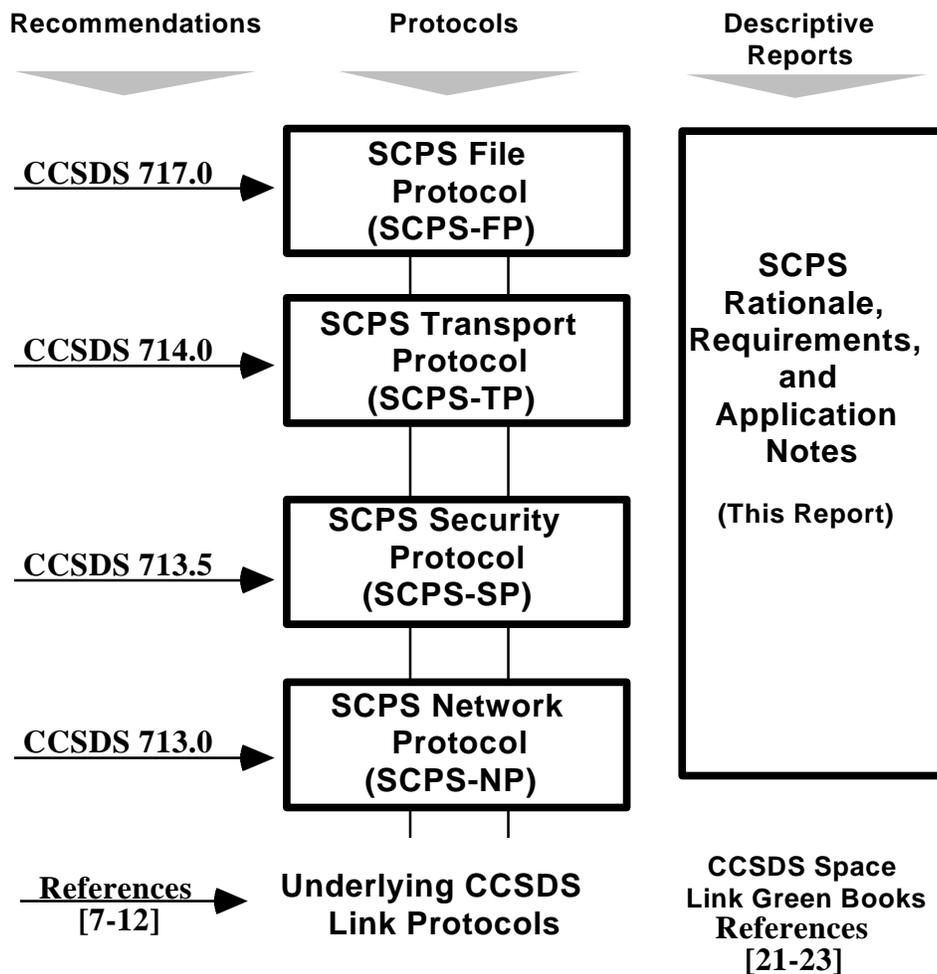


Figure 1-1: SCPS Layered Protocols

The SCPS protocols include:

- a file handling protocol (the SCPS File Protocol, or SCPS-FP), optimized towards the up-loading of spacecraft commands and software, and the downloading of collections of telemetry data;
- an underlying retransmission control protocol (the SCPS Transport Protocol, or SCPS-TP), optimized to provide reliable end-to-end delivery of spacecraft command and telemetry messages between computers that are communicating over a network containing one or more potentially unreliable space data transmission paths;
- a data protection mechanism (the SCPS Security Protocol, or SCPS-SP) that provides the end-to-end security and integrity of such message exchange;
- a scaleable networking protocol (the SCPS Network Protocol, or SCPS-NP) that supports both connectionless and connection-oriented routing of these messages through networks containing space data links.

This Report summarizes the rationale for development of the SCPS, and describes the services provided. The functional requirements for each of the protocols and summaries of the results of SCPS test and validation efforts are presented in appendices.

1.2 DOCUMENT STRUCTURE

This Report is organized as follows:

- Section 1 defines the purpose and scope of this Report and lists the definitions, conventions, and references used throughout the Report.
- Section 2 contains an overview of this Report. It presents the background of the SCPS project, and briefly describes the SCPS protocol stack and the individual protocols.
- Section 3 presents the rationale for the SCPS protocols and describes their relationship to earlier CCSDS and internet protocols.
- Sections 4 through 7 describe the four SCPS layers in bottom-up order: network, security, transport, and file handling, respectively. The features of each protocol are described, and the rationale for provision of those features is presented.
- Annex A contains a glossary of terms used in this Report.
- Annex B contains a list of acronyms used in this Report.
- Annex C provides a list of frequently asked questions (FAQ) and answers concerning SCPS.
- Annex D presents the Protocol Functional Requirements for the SCPS protocols.
- Annexes E and F contain the executive summaries of reports on test programs that have been conducted to validate the SCPS file and transport protocols and evaluate their performance.

1.3 DEFINITIONS

Concepts and terms that are necessary for overall understanding of this Report are presented here. Other terms and concepts are introduced where appropriate throughout this Report.

The SCPS Recommendations define communications protocols and services in the style established by the OSI Basic Reference Model, Reference [16]. This model describes communications services as being provided by layers of protocols (see Figure 1-2), each layer providing a service interface to users of the service in the layer above.

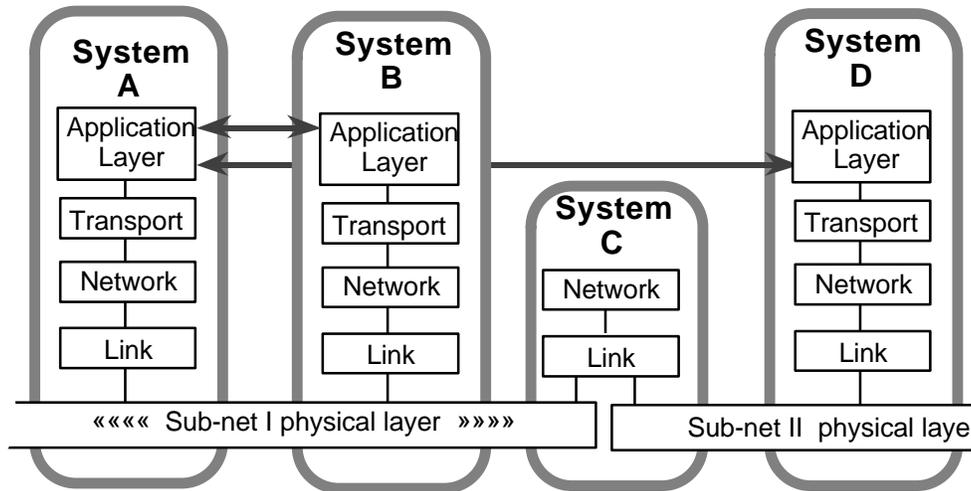


Figure 1-2: Layered Protocol Model

The concepts and terminology of the OSI Basic Reference Model are summarized in Reference [15]. Here, we show the layers that are common between the OSI model and the Internet; the Presentation and Session Layers located directly below the Application Layer in the OSI model are not shown. A few important aspects of this model are:

- a) In the uppermost layer, applications in one system interact with applications in other systems through the communications services of the lower layers.
- b) A layer comprises all of the peer communications entities in all of the communicating systems. In any one system, the entities of a given layer are called a subsystem. (e.g., the transport sub-system in System A might contain the Internet TCP and UDP protocol entities.)
- c) Sub-networks may be interconnected through relay systems, as illustrated by System C in Figure 1-2.
- d) Layers below the Transport Layer do not operate end-to-end. For example, when an application in System A sends data to an application in System D, The network entity in System A cannot know whether a data unit that reaches System C is actually relayed to System D. End-to-end functions must be provided at the Transport or Application Layers.

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 revisions, 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] Procedures Manual for the Consultative Committee for Space Data Systems, CCSDS A00.0-Y-6, May 1994, or later issue.
- [2] Exploratory Study: Recommended Development of Interoperable Data Communications Standards for Dual Use By US. Civil and Military Space Projects. SCPS D71.0-Y-1, Yellow Book, Issue-1, September 1993.
- [3] SCPS Network Protocol, Red Book, CCSDS SCPS-713.0-R-2, Issue 2, April, 1997, or later issue.
- [4] SCPS Security Protocol, Red Book, CCSDS SCPS-713.5-R-2, Issue 2, April, 1997, or later issue.
- [5] SCPS Transport Protocol, Red Book, CCSDS SCPS-714.0-R-2, Issue 2, April, 1997, or later issue.
- [6] SCPS File Protocol, Red Book, CCSDS SCPS-717.0-R-2, Issue 2, April, 1997, or later issue.
- [7] Packet Telemetry, Blue Book, CCSDS 102.0-B-3, Issue 3, November, 1992, or later issue.
- [8] Telemetry Channel Coding, Blue Book, CCSDS 101.0-B-3, May 1992, or later issue.
- [9] Advanced Orbiting Systems, Networks and Data Links, Blue Book, CCSDS 701.0-B-2, October 1989, or later issue.
- [10] Telecommand—Part 1 Channel Service, Blue Book, CCSDS 201.0-B-1, January 1987, or later issue.
- [11] Telecommand—Part 2 Data Routing Service, Blue Book, CCSDS 202.0-B-2, January 1987, or later issue.
- [12] Telecommand—Part 3 Data Management Service, Blue Book, CCSDS 203.0-B-1, January 1987, or later issue.
- [13] CCSDS Global Spacecraft Identification Field: Code Assignment Control Procedures, Blue Book, CCSDS 320.0-B-1, Issue 1, October, 1993, or later issue.
- [14] CCSDS Publications Manual, CCSDS A20.0-Y-1, May 1994, or later issue.
- [15] CCSDS Report: Terminology, Conventions, and Methodology, CCSDS 910.2-G-1, November 1994, or later issue.
- [16] Basic Reference Model for Open System Interconnection, ISO.

DRAFT REPORT: SCPS RATIONALE, REQUIREMENTS, AND APPLICATION NOTES

- [17] Internet Engineering Task Force, "Internet Protocol", (Postel, J.B., Ed.), Request for Comments (RFC) 791, (Postel, J.B., Ed.), September, 1981.
- [18] Internet Engineering Task Force, Postel, J.B.; and Reynolds, J.K., "File Transfer Protocol (FTP)," Request for Comments (RFC) 959, (Postel, J.B., and Reynolds, J.K.; Ed.), October 1985.
- [19] Internet Engineering Task Force, "Requirements for Internet Hosts -- Application and Support," Request for Comments (RFC) 1123, (Braden, R., Ed.), October 1989.
- [20] Internet Engineering Task Force, "TCP Extensions for High Performance," RFC 1323, (Jacobson, V.; Braden; R, Borman,.. D.; Ed.), May 1992.
- [21] "Telemetry: Concept and Rationale", Green Book, CCSDS 100.0-G-1, December 87, or later issue.
- [22] Advanced Orbiting Systems, Networks and Data Links: Summary of Concept Rationale, and Performance, Green Book, CCSDS 700.0-G-3, November 92, or later issue.
- [23] "Telecommand: Concept and Rationale", Green Book, CCSDS 200.0-G-6, January 87, or later issue.
- [24] Data Networks, Second Edition, by Dmitri Bertsekas and Robert Gallager. (1992, Prentice Hall).

2 OVERVIEW

2.1 BACKGROUND

Space missions have always had requirements for reliable, secure transfer of individual data units, or “messages,” and larger collections of data, or “files,” between space and ground. In the past these requirements have been met by a combination of custom-designed software and manual control by human operators, thus providing some functions of communications protocols intertwined with other functions of mission operations.

As shown in Figure 2-1, the space mission operations environment comprises the ground information infrastructure, with high-speed computing and communications capabilities, and the space information infrastructure, which presents a significantly different environment. In space there are relatively few end systems and networks, with performance far below that of ground nodes due to extreme mass, power, and volume constraints, together with the delay and expense of developing space-qualified hardware.

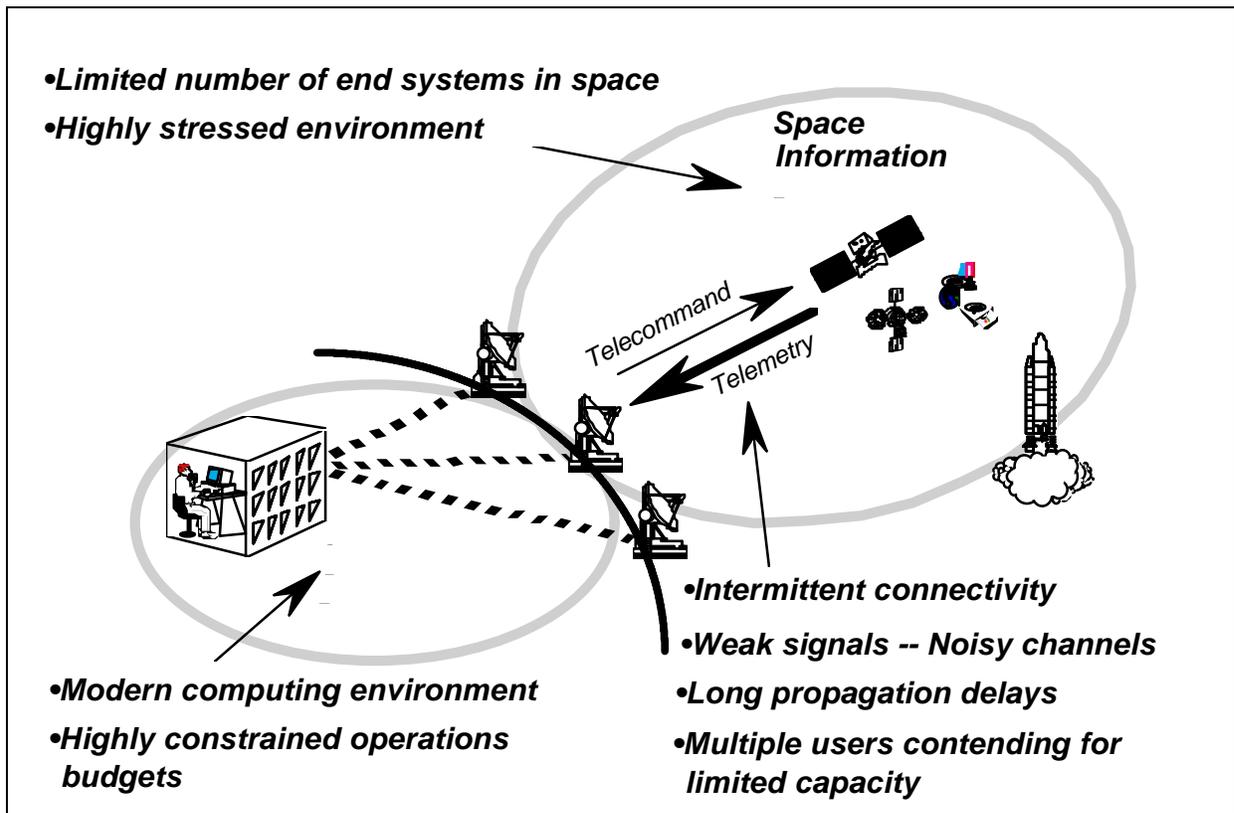


Figure 2-1: Extension of Internet into Space

Space communications is further complicated by the characteristics of the space/ground link. With rare exceptions, connectivity to a space vehicle is intermittent, with duty cycles typically below 10% due to limited visibility from ground stations and contention among missions for scarce contact time. Limited signal strength and noise make data loss through

corruption far more likely than in ground networks, and long propagation times cause terrestrial protocols to operate sluggishly or to fail outright.

SCPS protocols are designed to provide interoperability across the spectrum of space missions, and between space data systems and the broader ground network environment. They provide a set of options and protocol data unit formats that can be scaled to satisfy the communication needs of both complex and simple, resource-constrained missions.

The Internet protocol suite (e.g., TCP, UDP, IP, FTP) provides many functions needed for space communications, but these protocols are designed to meet environmental requirements that are significantly different from those encountered in communicating with a remote spacecraft. Although functionally equivalent to terrestrial networks, space communications networks often have performance and operational considerations that prevent direct use of existing commercial protocols. Today's internet protocols were developed for terrestrial networks and assume that connectivity is maintained, that data loss due to corruption is rare, that balanced bi-directional links are available, and that most data loss is due to congestion. Further, vendors of commercial communications products that implement these protocols use these assumptions to maximize performance and economy in this environment, making the treatment of retransmission, recovery, and time-outs inappropriate for space operations. For the large majority of current space programs, the space mission environment makes performance of these protocols unacceptable.

The SCPS supplement current space link and ground protocols with end-to-end protocols designed to bridge the space and ground environments (Figure 2-2).

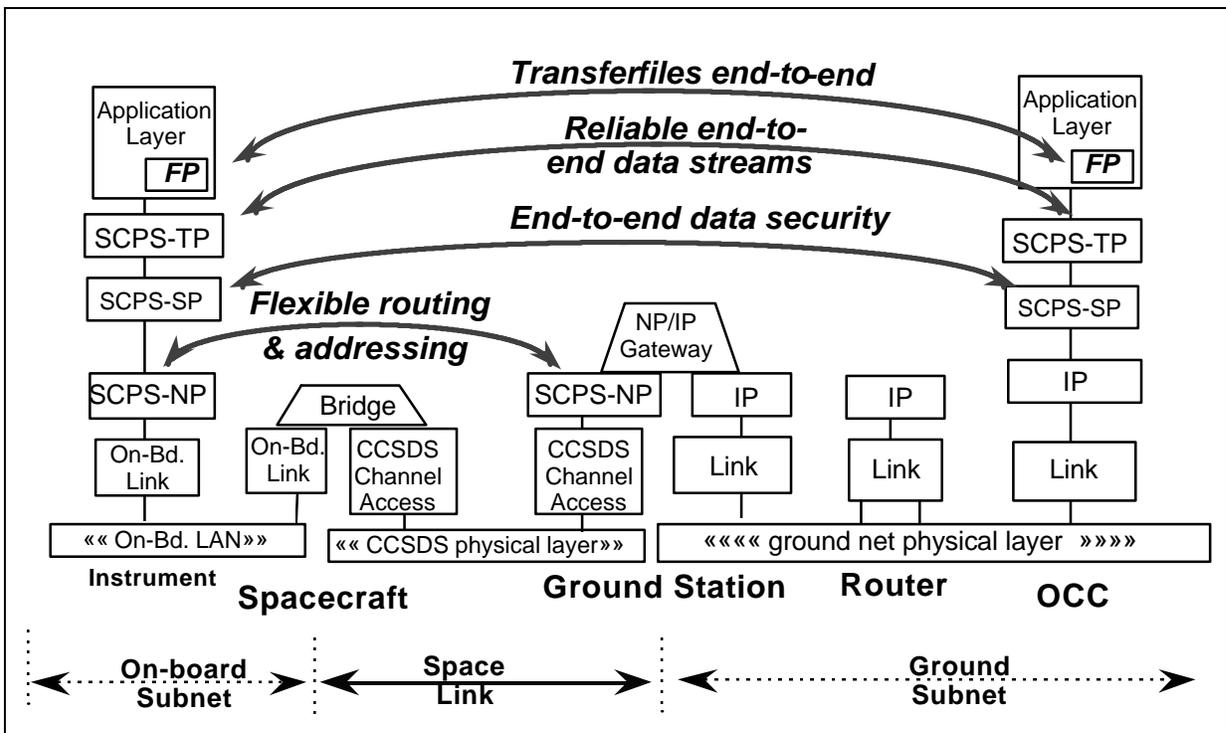


Figure 2-2: SCPS end-to-end Services

2.2 RATIONALE FOR SCPS

The principal goal of the SCPS effort was to lower lifecycle costs by reducing development and operations costs in space communications systems.

The SCPS program was initiated in response to several demands:

- a) A need for standard protocols to support reliable data transfer.
- b) The need to accommodate evolving multi-node mission configurations that require in-space network routing.
- c) The need to drastically reduce operations costs and thus maintain the ability to produce results from space missions in the face of decreasing funding.
- d) The need to provide compatibility and interoperability with the internet.

The SCPS are designed to meet these demands by increasing standardization and interoperability, both within and among CCSDS Agencies and other developers and operators of spacecraft.

SCPS augments previously developed CCSDS protocols by providing reliable stream or file transfer over CCSDS frames at the link layer and dynamic networking for those missions that need it.¹ CCSDS Packets remain as the CCSDS telemetry and telecommand source message format. CCSDS RF and modulation Recommendations are also unaffected. The SCPS were designed to operate over CCSDS space-ground links, although use of other link layer protocols is possible. Given the link characteristics and intermittent connectivity in space operations, a balance of upper-layer, confirmed, end-to-end services supported by links good enough to avoid excessive retransmission provides optimal performance.

The SCPS protocols provide flexibility and optional features that allow designers to tailor a communications protocol suite to meet the requirements and constraints of a mission, without extensive software development. Specific layers, or protocols within layers, can be included or omitted to create an optimal profile for the mission. Then each of the selected protocols can be adapted, if necessary, to meet specific mission requirements. Such adaptation is often made through compile-time options to tailor a standard protocol product for use in a particular environment. Some options are simply setup parameters that configure the run-time protocol entity to optimize performance or provide compatibility. Each of the options was designed to accommodate the differences in objectives, hardware, environment, or operations among space missions, without sacrificing the benefits of lower cost and risk provided by a set of coordinated standards.

¹At the time of publication, efforts were underway to enhance CCSDS link protocols to support multiplexing of SCPS Network Protocol and other network-level data units into CCSDS virtual channel frames.

2.3 APPLICABILITY OF SCPS

SCPS is aimed at a broad range of space missions including:

- a) Support for spacecraft in low-earth and geosynchronous orbits, as well as lunar and planetary spacecraft. The primary emphasis has been on support of missions at lunar distances or closer. SCPS network and security protocols are relatively immune to communications delay, and thus can support deep-space missions today. Additional capabilities for data transport and file handling in deep-space missions will be addressed in SCPS Phase 3, beginning in late 1997.
- b) Support of spacecraft with a range of on-board communication and on-board data handling resources, including those with limited on-board computer and memory resources, as well as those with multiple, high-capacity on-board computers with extensive data storage.
- c) Support of multi-node mission configurations, including:
 - 1) Cluster-like missions.
 - 2) Spacecraft constellations.
 - 3) Orbiter/lander planetary missions.

2.4 SCPS LAYERS/PROTOCOLS OVERVIEW

The SCPS protocols are based on widely-used terrestrial and space communications protocols, primarily internet protocols whose specifications are developed and maintained by the Internet Engineering Task Force (IETF). To meet the specific needs of space missions, modifications and extensions to these IETF protocols as well as concepts drawn from other protocols are incorporated in the design of the SCPS protocols. The modifications and extensions are discussed in SCPS FAQ questions 3.1, 3.2, and 3.3 in Annex C.

2.4.1. NETWORK LAYER—SCPS-NP

To avoid the high communications overhead of commercial internetworking protocols, the SCPS Network Protocol (SCPS-NP) was developed for in-space use. The SCPS-NP supports many different connectivity and routing environments. It supports different modes of operation - from highly managed to highly protocol-driven. Services support basic data transfer, local system support, and network diagnostics.

The SCPS-NP protocol header is designed to be scaleable, to allow a very wide range of in-space routing configurations to be supported via plug-in modules to the packet header. These modules provide subsets of the total spectrum of SCPS-NP capabilities. The SCPS-NP is based on the internet IP, and several proposed IP enhancements. Capability-driven header construction enables SCPS-NP to meet bit-efficiency requirements, and reduces resource requirements on missions with limited on-board data systems.

The services and protocol features of the SCPS-NP are further described in Section 4 of this Report. System engineering considerations are described in Section 8. The Specification for SCPS-NP is Reference [3].

2.4.2. SECURITY LAYER—SCPS-SP

Data protection services are based on capabilities similar to the ISO Network Layer Security Protocol (NLSF), as adapted from the Secure Data Network Systems (SDNS) "SP3" protocol. SCPS-SP provides options for data protection in space mission communications but with minimal communications overhead.

The services and protocol features of the SCPS-SP are further described in Section 5 of this Report. System engineering considerations are described in Section 8. The Specification for SCPS-SP is Reference [4].

2.4.3. TRANSPORT LAYER—SCPS-TP

To provide reliable end-to-end SCPS Transport Protocol (SCPS-TP) services, the Internet Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) were adapted to meet unique space mission requirements, using IETF-defined extensions and SCPS-defined modifications.

The services and protocol features of the SCPS-TP are further described in Section 6 of this Report. System engineering considerations are described in Section 8. The Specification for SCPS-TP is Reference [5].

2.4.4. APPLICATION LAYER: FILE HANDLING—SCPS-FP

For the SCPS File Protocol (SCPS-FP), both the Internet "File Transfer Protocol" (FTP) and the custom "Space Station Freedom File Transfer Protocol" were considered. FTP was selected as the basis for development of SCPS-FP, since this approach more closely meets the needs of a broad range of missions, and facilitates interoperability with FTP.

The services and protocol features of the SCPS-FP are further described in Section 7 of this Report. The Specification for SCPS-FP is Reference [6].

2.5 SCPS TESTING AND VALIDATION

The SCPS protocols have been tested to validate the specifications, and to evaluate their performance. Laboratory simulations were used throughout the development process to validate the algorithms used in the protocols. Prototype protocol implementations were tested in both simulated and actual space mission environments.

A satellite relay (bent-pipe) experiment to test SCPS-TP was carried out using a US Department of Defense satellite. SCPS-TP performed well, maintaining between 82% and 97% of maximum throughput (depending on packet size) at bit-error rates of up to 10^{-5} . As part of this test, the performance of SCPS-TP was compared to that of Transmission Control Protocol (TCP) using a similar configuration in the laboratory. SCPS-TP performance was equivalent to that of TCP at low bit error rates, and significantly better than TCP's at bit-error rates of 10^{-7} or greater. The executive summary of the SCPS Bent-Pipe Test Report is provided in Annex E.

The UK Defense Research Agency's Space Technology Research Vehicle (STRV) was utilized to exercise the SCPS protocols onboard an orbiting spacecraft. Several of the SCPS protocols (FP, TP, SP) were uploaded to the STRV and were tested between space and ground under actual flight conditions. Files were uploaded and downloaded between the ground and the STRV via the use of the SCPS File Protocol (SCPS-FP) and the SCPS Transport Protocol (SCPS-TP). SCPS-TP's ability to hold connections across short contact times, cope with high bit error rates on the space communications link, and its ability to provide high throughput were tested. The SCPS Security Protocol (SCPS-SP) was tested in conjunction with SCPS-TP and demonstrated that the SCPS-TP tests could be carried out in a secure environment. The executive summary of the SCPS-FP SCPS/STRV Flight Experiment (SSFE) Test Report is provided in Annex F.

2.6 SUMMARY

The primary goal of the SCPS effort was to reduce space mission operations costs by extending internet-like connectivity into space. The rationale for this approach is that both the data systems and the personnel (designers, operators, users) associated with space missions are already using internet protocols. The communications services that they need in space are very similar to those they have in ground networks. The easiest, lowest risk, and most direct way to achieve this goal was to adapt the protocols that are used on the ground. The SCPS, supported by link- and physical-layer protocols, provide the flexibility and performance needed for support of future space missions.

3 RATIONALE FOR SCPS

3.1 CHANGING ENVIRONMENT

Dramatic changes are occurring across many segments of the space community, driven by the combined forces of new technology, more demanding mission objectives, shrinking government budgets and renewed emphasis on developing commercial markets. Dimensions of the change include:

- a) a shift towards decentralization in mission strategy, with movement away from "a few expensive spacecraft launched relative infrequently" towards "many affordable spacecraft launched relatively often";
- b) increasing reliance on cooperation (both national and international) to achieve complex space mission objectives in ways that are affordable to individual organizations, coupled with an erosion of the traditional boundaries between the civil, military and commercial space sectors;
- c) consequent emphasis on reducing wasteful duplication of effort and improving mission effectiveness by sharing infrastructure via standards-induced interoperability;
- d) an overriding imperative to significantly reduce mission operations costs via increased automation, decreased man-in-the-loop and 'specialist' operations, and fewer space link sessions, and an increased reliance on commercially-derived capabilities which are provided by the private sector.

These changes in the mission environment have forced a large scale re-thinking of the way in which space missions are executed. Mission operations costs must be drastically reduced by eliminating labor-intensive activities, replacing them with highly automated approaches. Mission planners are increasingly emphasizing the important role that standardization plays in achieving significant reductions in system development and operations costs.

3.2 OPERATIONAL CONSTRAINTS

Ideally, the challenges described above would be met through use of off-the-shelf technology that has been proven in ground-based systems. Unfortunately, space missions must be carried out under conditions that vary significantly from those in ground data systems. The operational constraints encountered in space communications include:

- a) Round-trip delays much greater than those seen in ground networks.
- b) Noise characteristics on space links that, despite sophisticated error correction codes, produce more frequent data loss than on ground links.
- c) Variation in the format and performance characteristics of the space links used in space missions. Within the CCSDS link protocols, different levels of error protection are available.
- c) Intermittent connectivity, as a result of orbital position, earth rotation, or availability of ground station support.

- d) Changes in the routing path from contact to contact, because of use of multiple ground stations or changes of the relative positions of multiple spacecraft.
- e) Low forward bandwidth, or, more generally, asymmetry between the forward and return bandwidth. Virtually all missions, other than those dedicated to 'bent-pipe' communications service, require a much higher return data flow, compared with the forward data rate needed for commanding and maintenance. This asymmetry has effects on the features of protocols that support end-to-end communications, as noted in Section 3.3.2.

3.3 EXISTING PROTOCOLS

3.3.1. CCSDS SPACE LINK PROTOCOLS

3.3.1.1 Current CCSDS Capabilities

In the area of standards for space data communications and flight operations, the community is fortunate to be able to build on the foundations laid by CCSDS. The first wave of CCSDS Recommendations, which focus primarily on the data link interconnecting the spacecraft with its ground support system, introduced two sweeping new capabilities:

- a) asynchronous packetized data transfer, which unshackles the internal operations of spacecraft systems from being in lock-step with the time-division multiplexed space/ground data communications process. This allows more efficient use of link capacity by adaptive and event-driven telemetry and telecommand applications;
- b) high performance channel coding, which has virtually eliminated the space link as a source of undetected bit errors and has thus made a significant stride towards supporting both data compression and true computer-to-computer data exchange between spacecraft and their supporting ground systems.

3.3.1.2 Additional Requirements

Despite these advancements, spacecraft and their ground systems are still unable to conduct automated computer-to-computer dialog of the kind that is routinely supported at very low cost on the Internet. The problem is rooted both in the current low-level of spacecraft automation, and in the absence of end-to-end data communications capabilities that perform acceptably in the space mission environment. In particular:

- a) The space link is only one component of the end-to-end data path between the user and a remote space investigation. There is currently no space-proven standard mechanism available to ensure that the end-to-end data transfer is fully reliable.
- b) As onboard computers become increasingly capable and onboard storage shifts from tape recorders to solid state memories, more and more telecommand and telemetry applications will become file oriented. Currently, there is no space-proven standard mechanism available to support end-to-end file transfer.
- c) The current CCSDS telemetry and telecommand capabilities rely on relatively simple spacecraft configurations with static routing relationships between end systems in space and on the ground. As space systems become more diverse there will be new

requirements to route data dynamically through changing in-space network topologies. Currently, there is no space-proven standard mechanism available to efficiently support such connectionless data routing. This problem is particularly acute in fleets or constellations of small spacecraft or lander vehicles with in-space crosslinks. It is also significant even for single spacecraft supported through several ground terminals or with several data sinks or sources on the ground.

- d) Space systems have traditionally tended to rely on their uniqueness to deter unauthorized access. As Internet connectivity becomes ubiquitous and space systems become integrated with the global communications infrastructure, there will be an increasing danger of malicious intrusion or unauthorized access to space vehicles and the sensitive information flowing within them. There is currently no space-proven standard mechanism available to ensure end-to-end space data protection.

3.3.2 INTERNET PROTOCOLS

Most space missions have had to deal with the problems of performing reliable, secure file transfers between space and ground, and have expended considerable resources either designing customized protocols or using the (expensive) reasoning power of human operators to provide the needed functions. Space missions have always had requirements for upper-layer functions; what they have lacked is standard protocol solutions. There are networking protocols that provide end-to-end capabilities, but they have two shortcomings in supporting space missions:

- a) The protocols are not designed to operate under the conditions encountered in space missions.
- b) The designers of commercial products that implement these protocols make assumptions that are reasonable in ground networks, but are significantly different from those encountered in space missions, leading to poor use of bandwidth and contact time, and loss of data.

The most widely used protocols today are the internet protocols. These are usually referred to as TCP/IP, but, in fact, comprise more than fifty internet standards. This communications baseline is robust and flexible, as the result of hundreds of thousands of engineering hours and years of use and testing. The SCPS provide modifications and extensions to only a few of these internet protocols, in order to meet the special requirements of space communication. Caution, economy, and compatibility are inherent in this approach.

3.3.2.1 File Handling

The SCPS-FP is derived from the Internet File Transfer Protocol (FTP). Like FTP, SCPS-FP uses two transport connections between host systems—a control connection to exchange control information, and a data transfer connection to move file data. Data is transferred from a storage device in the sending host to a storage device in the receiving host. FTP provides much of the functionality required for space mission operations, but does not address the resource restrictions of the space operations environment.

Both contact time and bandwidth are scarce resources in space operations. SCPS-FP operates under these constraints by providing enhanced error recovery and restart capabilities. Thus interruptions in file transfer can be restarted from the point of interruption instead of starting over, as would be necessary with FTP. SCPS-FP also provides the capability to read or update part of a file on a remote system rather than the entire file. This avoids the transfer of a large amount of data when only a small part of the file is affected. Other SCPS-FP extensions to FTP provide integrity checks to recover from errors in file transfer or update operations. Finally, to conserve bandwidth and contact time, SCPS-FP suppresses text messages between hosts involved in file operations. Further details are provided in Section 7 of this report.

3.3.2.2 Transport

Large propagation delays, limited bandwidth, losses due to errors, asymmetric link capacities, and intermittent connectivity all conspire to limit TCP's performance over space links. In many cases, TCP can cope with a subset of these environmental obstacles, although the performance achieved is far from optimal. The SCPS-TP improves performance in the space environment through a set of TCP extensions. Some of these extensions are changes to the TCP specification, while others are implementation details that do not affect interoperability.

To some degree, forward error correction can compensate for such errors, but the space link still is rarely as clean as those of terrestrial networks. TCP handles packet loss by retransmitting lost segments; however, TCP assumes the source of all packet loss is network congestion. Consequently, in response to packet loss, TCP invokes congestion control, reducing its transmission rate. This response is inappropriate when data loss is due to corruption rather than congestion, as is often the case on space links. TCP's congestion control algorithm works well in dealing with congestion-induced loss, but unnecessarily lowers throughput on uncongested, noisy links.

SCPS-TP provides the means to distinguish among the three possible causes of data loss—congestion, corruption, or link outage—and to invoke appropriate algorithms to deal with each of these. Further details are provided in section 6 of this report.

3.3.2.3 Security

State-of-the-art security protocols include the Secure Data Network (SDNS) "SP3" protocol, the ISO Network Layer Security Protocol (NLSP), the Integrated Network Layer Security Protocol (I-NLSP), and the Internet Engineering Task Force's (IETF) Internet Protocol Security (IPSEC) Encapsulating Security Payload (ESP) and Authentication Header (AH) protocols. All of these security protocols provide data confidentiality, data integrity, authorization, and access control, but they have far greater overhead than is acceptable over space links. The SCPS-SP has been designed to provide the security features needed for space operations with minimal overhead. Further details are provided in Section 5 of this report.

3.3.2.4 Network

Space networks tend to be more bandwidth-constrained (especially on their forward links) than terrestrial networks. The combined constraints of power and weight make bandwidth a scarce commodity. The Internetwork Protocol (IP) has a fixed minimum header size that includes protocol information related to features that have limited or no use in space systems. Further, some space networks require capabilities that are not supported by IP, and their addition would further increase the IP header size (through the use of IP options).

The SCPS-NP addresses the problems of bandwidth constraint by providing a scaleable header, containing only the header elements required by a particular packet. In addition, the SCPS-NP supports address translation, so that networks of limited scope are not required to carry large addresses. However, the SCPS-NP provides the ability to carry addresses appropriate to terrestrial networks, to accommodate very large space-based networks. Further details are provided in Section 4 of this report.

3.4 SUMMARY

In summary, while there is a clear need to extend Internet-like services into space, there is an equally clear need to engineer the protocol solutions to accommodate the highly stressed nature of the space environment. In particular:

- An efficient network layer capability is needed to provide many IP-like features without the overhead of IP itself and which also handles unique space routing problems.
- Efficient and powerful end-to-end data protection mechanisms are needed; these do not currently exist in the commercial marketplace.
- A Transport layer is needed to provide TCP- and UDP-like services that accommodate the non-terrestrial environment of computationally-constrained spacecraft end systems, long propagation delays, frequent interruptions in service, asymmetric channels and the presence of corruption rather than congestion as a principal source of data loss.
- An FTP-like file transfer service is needed which accommodates the need to provide basic file transfer services within an environment of constrained spacecraft end systems and frequently- interrupted communications.

4 SCPS NETWORK PROTOCOL (SCPS-NP)

4.1 SCPS-NP OVERVIEW

The SCPS Network Protocol operates at Layer 3 of the OSI Basic Reference Model. Its primary goal is to route data from the source of the data to its ultimate destination with the user's requested quality of service. The SCPS Network Protocol operates at *intermediate systems* that decide how to forward packets to their destinations. This is shown in Figure 4-1.

The SCPS-NP is a new protocol. It is not a subset of the Internet Protocol (IP) [Ref. 17], although it draws on concepts and technology from IP, and shares some IP numbering in its service interface. The SCPS-NP also draws on concepts and technology from the Path Service of the CCSDS Recommendation for Advanced orbiting Systems [Ref. 9], but is neither a subset nor a superset of that protocol. The Specification of SCPS-NP is provided in a CCSDS Recommendation, Reference [3].

The drivers for generating a new protocol were to provide high bit-efficiency (primarily compared to IP), and to provide capability for user-specified qualities of service not provided by either IP or the existing CCSDS capabilities.

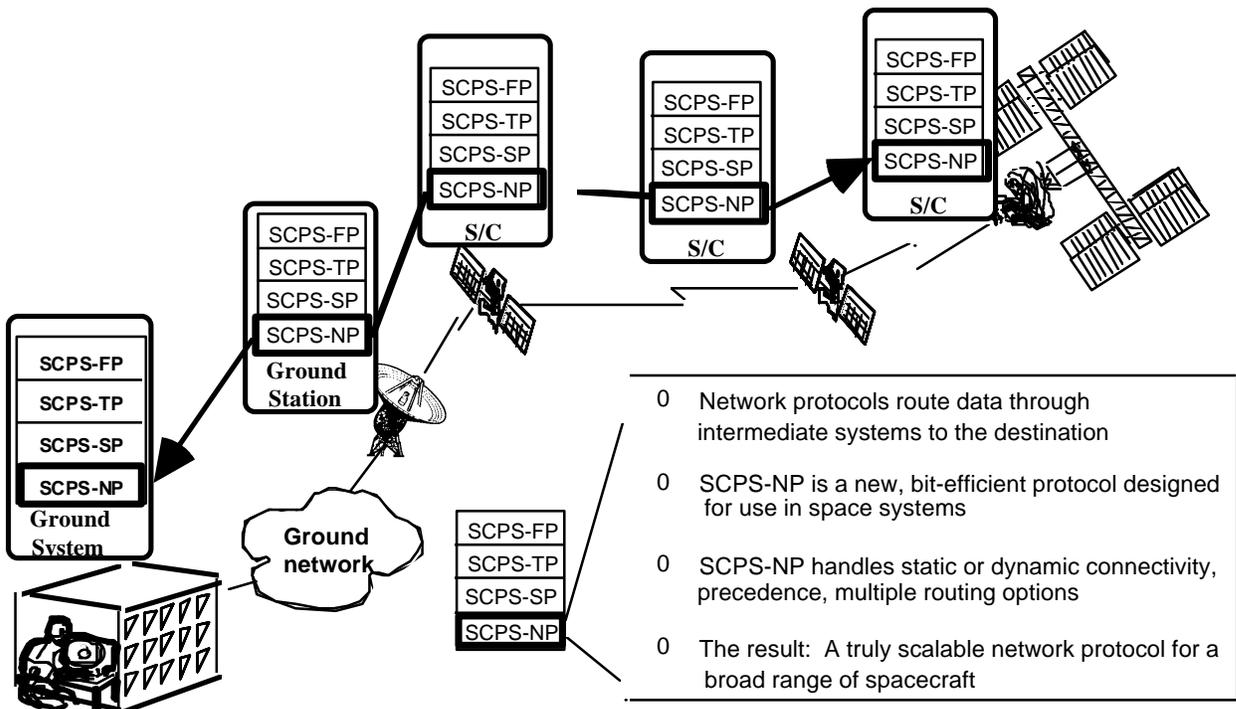


Figure 4-1: SCPS Network Protocol (SCPS-NP)

Section 4.2 presents some terminology useful to the subsequent text, then Section 4.3 presents and discusses the requirements that were allocated to the network layer. Section 4.4 examines how the IP might be used to meet those requirements and identifies some of the shortcomings of such an approach. Section 4.5 describes the capabilities within the SCPS-

NP that address the requirements. Section 4.6 discusses configuration alternatives for use of SCPS-NP in IP-based ground networks.

4.2 TERMINOLOGY

In this section, we present terms that are appropriate to the subsequent discussion and are either not widely known or are prone to misinterpretation.

A *host* is network-addressable system that may send or receive network-layer packets, but does not forward packets.

A *router* is a network-addressable system that may send, receive, or forward network-layer packets.

A *gateway* is a network-addressable system that terminates a protocol at a given layer and invokes similar services at the same layer of an adjacent network.

An *intermediate system* is a router or a gateway.

A *connection*, in communication terms, is a term that describes information that is named, persistent, and shared across the systems supporting the communication. Data sent via that connection make use of the shared state, thus gaining bit-efficiency and possibly processing advantages. One implication of the use of a connection is that all data flowing on the connection is treated in the same manner, as specified by the state information that defines the connection. Rather than carrying the state information itself, the data is accompanied by an identifier that is used to reference the state information. This state information is typically the source and destination, but may include information such as precedence. A "managed connection" is one in which the shared state is distributed via network management mechanisms (outside the scope of the SCPS-NP).

Some refer to a *datagram* as a unit of data transmission for connectionless networks and a *packet* as a unit of transmission for connection-oriented networks. We make no such distinction here - the terms are used interchangeably to represent a variable-length, octet-aligned protocol data unit. Neither the term *packet* nor the term *datagram* imply any particular layer of the OSI reference model when used in this document.

The term *flood routing* describes a routing technique that is used to improve the probability of receipt of important packets. In flood routing, a packet is replicated and transmitted to all adjacent nodes (routers, gateways, or hosts). The adjacent nodes that are routers or gateways then replicate the packet and repeat the process. This technique is typically applied in networks with rich connectivity (meaning that each node is connected to several other nodes). Packet identification techniques are used to prevent a node replicating a packet more than once. Flood routing generates a substantial amount of traffic, so it is used sparingly. However, its use improves the probability that all nodes in the network will receive at least one copy of the packet (the packet identification techniques ensure that *only* one copy of a particular packet will be delivered to upper layer protocols). Appropriate uses are for very high priority traffic and for routing updates that should be delivered to all systems in the network.

4.3 SCPS-NP REQUIREMENTS

This section summarizes the technical requirements that have been allocated to the network layer, and provides some discussion of how those requirements relate to SCPS communications environments. Prospective users and network designers should consider which of these requirements apply to their operational environments, in order to make appropriate configuration decisions based on information presented later in this section.

The key technical requirements that were allocated to the network layer are summarized below. The SCPS Network protocol must

- Route data from source to destination
- Provide efficient operation in constrained-bandwidth environments
- Provide precedence- (priority-) based data handling
- Provide packet lifetime control
- Provide selectable routing treatments
- Provide signaling of network conditions to upper layer protocols

The ability to route data from source to destination is characteristic of essentially all protocols that operate at the network layer of the OSI Basic Reference Model. The SCPS Network Protocol was based on requirements derived from several types of network architectures ranging from simple to complex. The simplest networks involve a single link and dedicated end systems at either end of that link. Other networks involve a single destination end system (satellite onboard computer) communicating through one or more ground stations over a ground network to an operations center consisting of several end systems. Some spacecraft may have onboard networks. The more complex networks that served as sources of requirements involve networks with changing topologies, such as those found in satellite constellations or in mobile radio networks. In these topologies, end systems may communicate with other mobile end systems or may communicate through the mobile network to the ground-based (wired) network.

Common to all of the prospective environments is that bandwidth may be constrained, either unidirectionally or bidirectionally. This constraint results in a requirement to operate with high bit-efficiency. Bit-efficiency quantifies the fraction of transmitted bits that are user data. Improving bit-efficiency may be accomplished in two ways: by increasing the amount of user data per unit of protocol control information (i.e., header information), or by decreasing the amount of protocol control information per unit of user data. The first approach, making packets longer, is simple, but does not work well in environments that are prone to bit-errors. It also does not work well when the user's data does not lend itself to aggregation. The second approach, reducing protocol header overhead, is the approach used throughout the SCPS Network Protocol design. Several requirements derive from the need to operate in bandwidth-constrained environments: multicasting, support for managed-connections, and precedence-based data handling all address bandwidth-related constraints and are described in subsequent paragraphs.

Multicasting is a technique for improving network-wide bit-efficiency. The technique of multicasting allows addressing of data to a group of destination systems. Rather than sending a unique copy of the data to each remote system, data are sent to the group address, and intermediate systems replicate the multicast packet only as necessary in order to reach all of the destination systems in the multicast group.

Managed-connections can enhance bit-efficiency in networks that can be characterized as having a few source-destination pairs that account for most of the network's traffic. For these flows, the source and destination addresses can be replaced by an identifier for the managed connection. In the SCPS-NP, this identifier is called a *Path address*.

Precedence improves operation in bandwidth-constrained environments in two ways. First, it controls the order of service, which reduces queuing delay and variation in queuing delay for high precedence traffic. Second, precedence controls the order of packet discarding when congestion occurs, to ensure that if packet discarding is required, low precedence packets are discarded before higher precedence packets.

Packet lifetime control provides protection against transient routing loops. A transient routing loop is formed when routing tables in the network are not synchronized. This condition can occur as a result of using certain routing protocols, such as Shortest Path First Reference [24]. While a routing loop is in existence, the links forming the loop may become progressively more congested. Packet lifetime control ensures that data packets do not remain in the network indefinitely, as they are discarded once they have exceeded their "lifetime." This, combined with either automated or manual means to update the routing tables provides control over routing loops.

The requirement for selectable routing treatments provides the ability to switch between "normal" routing and other routing treatments, such as flood routing. This is of use in the more complex network topologies that involve relatively rich connectivity, such as satellite constellations with cross links or some mobile radio networks. The ability to flood route packets in these networks can improve the probability of receipt and reduce the propagation time of the flood routed packet through the network.

Signaling of network conditions to upper-layer protocols is required to allow those protocols to become aware of and to adapt to changing conditions within the network. Signals that may be passed to the upper-layer protocols include indications of network congestion, network corruption, and link outages. This requires the network to identify these conditions at points in the network that may be remote from the end systems that host the upper-layer protocols, and to propagate network-internal signals to the affected end systems for delivery to the upper-layer protocols.

4.4 SHORTCOMINGS OF USING IP IN SPACE NETWORKS

The IP is a highly capable, widely distributed protocol. It is an appropriate protocol for many environments, and may be appropriate for use in some SCPS environments. Due to its broad commercial support, if it will meet the requirements of a particular mission, it should be given serious consideration.

Table 4-1 lists the requirements presented in Section 4.3 and identifies whether IP meets those requirements. Clearly, IP can route data from its source to its destination, although as with the SCPS Network Protocol, the choice of specific routing protocol is dependent on the local networking environment.

Table 4-1. Support of SCPS Network Requirements by IP

Requirement	Support in IP?
Route from source to destination	Yes
Support for constrained bandwidth	No
Multicasting	Yes
Managed-connection operation	No
Precedence- (priority-)based handling	Partial
Selectable routing treatment	No
Packet lifetime control	Yes
Signaling to support upper-layer processing and network control	Partial

In general, IP does not provide any explicit support for operating in constrained-bandwidth environments. IP headers are a minimum of 20 octets in length, and may be made longer with the addition of options. IP provides support for multicasting, but has no mechanism for shortening its headers by using managed connections.

The IP header contains a field to carry eight levels of precedence. However, commercial equipment typically does not make any use of the field. In particular, high precedence packets would not benefit from any reduced probability of discard in congested routers, nor would they receive any reduced queuing delays in routers.

There is no concept of flood routing in IP. While an IP option could be defined to signal flood routing, there is no routing support for it in commercially-available implementations.

The capabilities in IP for packet lifetime control are adequate for most environments.

With respect to signaling of network conditions, some IP implementations provide partial support. The Internet Control Message Protocol (ICMP), the companion protocol to IP that handles such signaling, has the ability to generate congestion signals. However, the use of this signal has been deprecated due to the inability of routers to control the rate at which the congestion signals are generated (this problem may have been solved with the advent of Random Early Detection (RED), but RED is not widely deployed nor is its Explicit Congestion Notification (ECN) option). There is no signaling provided to indicate loss due to corruption nor to link outage.

4.5 CAPABILITIES OF THE SCPS NETWORK PROTOCOL

The SCPS Network Protocol provides support for all of the requirements identified in Section 4.3. The protocol is designed in such a way that unnecessary header elements are not incorporated into the header. This design decision increases the processing to format and parse SCPS-NP headers in favor of reducing the number of bits that are transmitted.

Table 4-2 reprises the requirements presented in Section 4.3 and identifies the support for those requirements within the SCPS-NP. The following paragraphs describe the capabilities of the SCPS-NP and how they meet the requirements.

Table 4-2: Support of SCPS Network Requirements by SCPS-NP

Requirement	Support in SCPS-NP
Route from source to destination	Yes
Support for constrained bandwidth	Short, variable length headers
Multicasting	Yes
Managed-connection operation	Path addressing
Precedence- (priority-)based handling	Yes
Selectable routing treatment	"Normal" and Flood Routing
Packet lifetime control	Hop Count and Time stamp
Signaling to support upper-layer processing and network control	Separate signals for congestion, corruption, and link outage

Network protocols typically route data from source to destination by selecting the "next hop" router based on the destination address. There are many methods by which the next hop router is selected. The SCPS-NP selects its next hop router by means of routing tables, which may be statically or dynamically configured. Routing tables that are statically configured are typically maintained either with network management or by distributing the tables in files. Some network configurations benefit from the use of routing protocols to maintain the routing tables. The routing protocols are not part of the SCPS-NP, but interact with the SCPS-NP routing tables.

The SCPS-NP is designed for constrained-bandwidth operation. The protocol has only a few elements that are present in every packet header: a version number, the packet length, the transport protocol identifier, and a control field. The transport protocol identifier indicates the network user (e.g., the SCPS-TP's TCP, UDP, or Compressed TCP, or the SCPS-SP) to which the packet's data should be delivered. The control field is a variable-length bit-field that signals what protocol header elements appear in the remainder of the header. These optionally-appearing header elements include both source and destination addresses, fields for precedence and routing requirements, fields to support packet lifetime control, and a header checksum for header error detection.

The SCPS-NP may perform address translation to improve bit-efficiency. Typically, the protocols that use SCPS-NP operate using IP addresses. This is especially important to consider if the SCPS-NP protocol operates only in the space (or wireless) segment of the network and protocol translation rather than encapsulation is performed. Section 4.7

discusses this topic further. However, IP addresses are four octets in length, and there are two of them. These can be carried without translation in the SCPS-NP header, or may be translated into more bit-efficient representations. The address formats are described in the SCPS-NP specification. An IP source-destination pair may be translated into three alternate versions: an Extended Path Address, which represents the two addresses with a single four-octet address; a pair of Basic End System Addresses, which represents the two four-octet addresses with a pair of corresponding one-octet addresses; or a Basic Path Address, which represents the pair of four-octet addresses with a single one-octet address. The use of a single address to represent an address pair is what is meant by a *managed connection*. The SCPS-NP decides whether to translate addresses based on address translation tables that are configured statically. These translation tables are identical throughout the network.

The SCPS-NP supports multicasting, and identifies multicast (group) addresses based on the address format. Either end-system addresses or path addresses may signify multicast address groups. (Refer to the SCPS-NP specification for the details of multicast address formats.) Currently in the SCPS-NP, the systems that belong to multicast groups are defined statically.

The SCPS-NP carries precedence in its Basic Quality of Service header field. When packets are queued for transmission within the SCPS-NP, the precedence field controls the order of transmission (higher-precedence packets are transmitted first), and the order of packet discard in the event of congestion (lower-precedence packets are discarded first).

The SCPS-NP supports selection of routing treatments as another element within the Basic Quality of Service header field. Signaling of four different routing treatments is supported by the protocol, with two of those being defined in the current version of the protocol. The two routing treatments that are defined are "normal" and flood routing, and are described in Section 4.3, above.

The SCPS-NP supports two forms of packet lifetime control. The first uses a hop-count, which is initialized at the packet's source and decremented every time the packet is routed. If it reaches zero before the packet reaches its destination, the packet is discarded. (This prevents packets that are caught in routing loops from remaining in the network indefinitely.) All networks that have the possibility of routing loops can make use of the hop count capability. The second form of packet lifetime control is based on a source time stamp, which is carried in the SCPS-NP header and indicates the time at which the packet was sent. A system that receives the SCPS-NP packet checks the time stamp to determine whether the packet is too old to be forwarded. This decision is made based on a Maximum Packet Lifetime configuration parameter. This form of packet lifetime control depends on having synchronized system clocks on all of the systems that host the SCPS-NP. It offers a bit-efficiency advantage in some cases, in that the source time stamp may be the same one used by the transport layer for round-trip timing. In this version of the SCPS Network Protocol release, clock sources are not assumed to be synchronized throughout the network. When packet lifetime control is required, the hop count parameter is used by default. Source time stamps submitted by the transport protocol will be carried in addition to the hop count field.

The SCPS-NP provides the SCPS Control Message Protocol (SCMP) to accomplish necessary signaling between SCPS-NP entities. It supports essentially similar error and query services as are found in the Internet Control Message Protocol (ICMP), but with

additions for corruption-experienced and link-outage signaling. (ICMP already has signaling to report congestion - it is called the "Source Quench" message.)

4.6 CONFIGURATION ALTERNATIVES

There are two main alternatives for configuring the SCPS Network Protocol to operate in IP-based ground networks. Each alternative has its own advantages, which are discussed in this subsection. The two alternatives may be referred to as "encapsulation" and "translation".

4.6.1 ENCAPSULATION

Figure 4-2 illustrates the basic concept of the encapsulation approach: SCPS-NP packets are formed at the data source, and routed through the ground network by *encapsulating* them in IP packets or UDP/IP packets. On the right half of the drawing, which represents the typical ground network, IP packets are used to carry SCPS-NP packets that, in turn, carry TP packets and user data. (In this drawing, link headers are not shown, and the SCPS-SP is assumed to be not in use.) The center box in the figure represents a router at the point where the wired network meets the wireless (space) network. In this router, the IP header is removed (for space-bound packets) or added (for packets coming from the space-based portion of the network). On the left half of the drawing, which represents the wireless portion of the network, packets do not carry the IP headers, reducing header overhead.

The primary advantage of the encapsulating approach is that the signaling capabilities of the SCPS Control Message Protocol are preserved end-to-end throughout the network. At routers and end systems that support the SCPS-NP, signals indicating congestion, corruption, and link-outage can be generated (assuming that those conditions can be detected in the local system). Note, however, that congestion or corruption loss or link outages that occur *within* the IP network are not signaled by the NP. Until the techniques of RED with Explicit Congestion Notification are widely deployed throughout the Internet, congestion signaling will not be available. Corruption is not a significant problem within the Internet, so no signaling for corruption is currently necessary. Likewise, link outage is not currently a significant problem within the wired portions of the Internet.

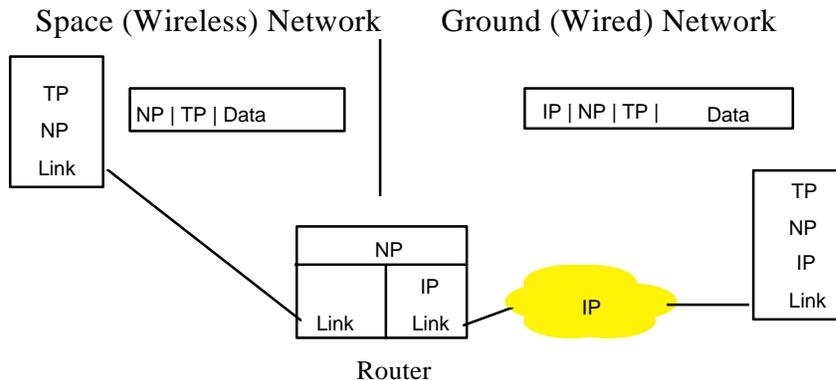


Figure 4-2. Encapsulating Network Configuration

The primary advantage of the encapsulating approach is that the signaling capabilities of the SCPS Control Message Protocol are preserved end-to-end throughout the network. At routers and end systems that support the SCPS-NP, signals indicating congestion, corruption, and link-outage can be generated (assuming that those conditions can be detected in the local system). Note, however, that congestion or corruption loss or link outages that occur *within* the IP network are not signaled by the NP. Until the techniques of RED with Explicit Congestion Notification are widely deployed throughout the Internet, congestion signaling will not be available. Corruption is not a significant problem within the Internet, so no signaling for corruption is currently necessary. Likewise, link outage is not currently a significant problem within the wired portions of the Internet.

4.6.2 TRANSLATION

Figure 4-3 illustrates the translation approach to routing through the ground network. In this approach, the ground-based system on the right supports IP, but not the SCPS-NP. At the gateway in the center of the figure, for ground-bound packets the information from the NP headers are used to form IP headers, which replace the NP headers. Similarly, space-bound packets have the IP headers removed and replaced with NP headers containing similar information.

The main advantage of this approach is that the ground-based system on the right may use either a commercial implementation of TCP (at a loss of the SCPS enhancements) or a SCPS-enhanced TCP implementation over regular IP. The SCPS Control Message Protocol's congestion signal can be propagated to the ground-based system via the Internet Control Message Protocol (although many commercial TCP implementations do not respond to it or respond incorrectly).

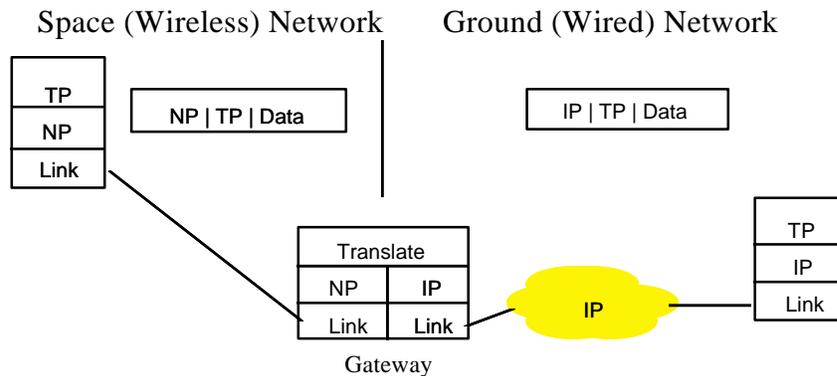


Figure 4-3. Translating Network Configuration

5 SCPS SECURITY PROTOCOL (SCPS-SP)

5.1 SCPS-SP OVERVIEW

The SCPS Security Protocol (SCPS-SP) provides *end-to-end* security services with very low bit overhead. The SCPS security protocol is an optional SCPS protocol and is algorithmic independent. No confidentiality (e.g., encryption) or integrity algorithms are specified as part of the protocol. The SP is physically located between the transport and network layers, as shown in Figure 5-1 and is designed to operate at the upper portion of layer 3 (network). In this manner, it can be implemented in both end-systems as well as in intermediate systems. It can operate over SCPS-NP, or other connectionless network protocols such as the IP. It can also operate over existing CCSDS protocols via a convergence layer.

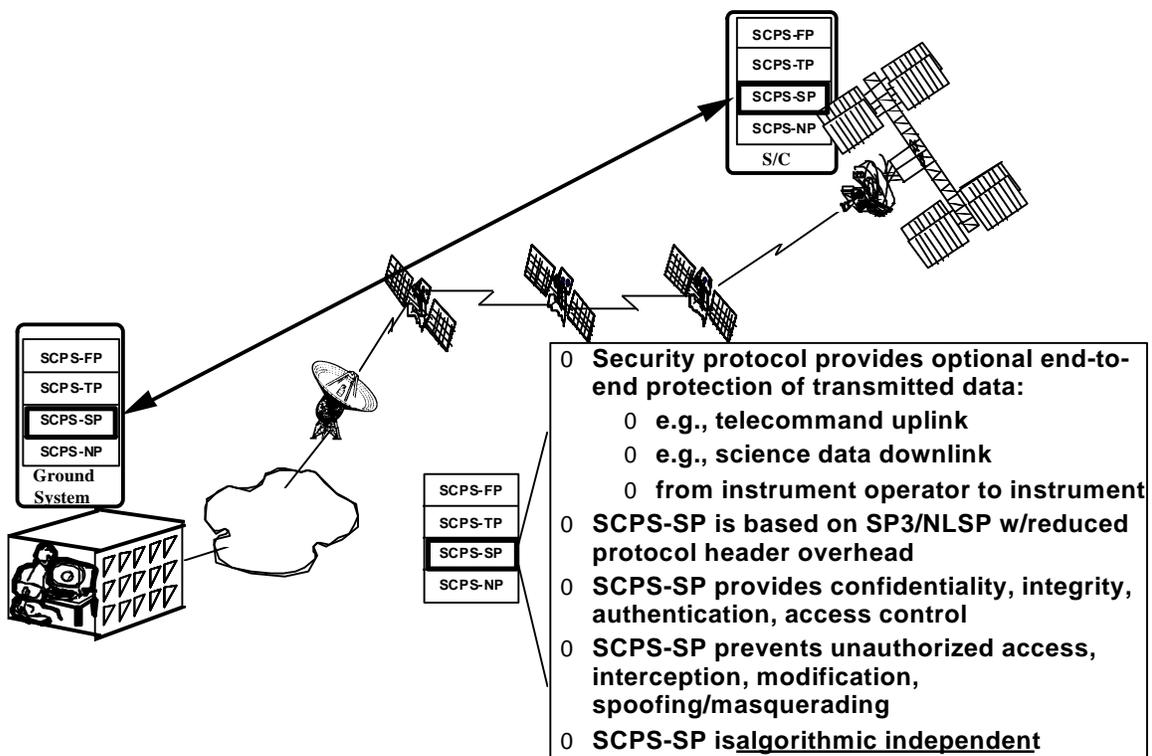


Figure 5-1: SCPS Security Layer

SCPS-SP provides security services on an *end-to-end* basis - from the source of the transmitted data to its destination. Any non-security-related intermediate systems (e.g., routers, gateways, control centers) will not have access to the data and will be prevented from viewing or modifying the data unless explicitly authorized. The communication end-points are implementation-specific and are defined by the implementing system.

For example, an instrument control center could be one endpoint and an instrument on-board a spacecraft could be another endpoint. In this manner, any data being relayed through a spacecraft control center or a ground station could not be viewed or altered.

However, another implementation might use an instrument control center as one endpoint with the ground station as the other endpoint. In this way, the data would be protected throughout the ground network but would then be unprotected within the ground station. Link layer security could be used between the ground station and the spacecraft to provide protection across the space link.

In order to provide end-to-end security services, SCPS-SP must allow the headers from the layers below it (e.g., network and media access addresses) to remain readable to allow an existing, unmodified network infrastructure to route the secured PDUs. Because of this, SCPS-SP does not provide protection against traffic analysis and enciphered data may be intercepted. Traffic analysis and low probability of interception may only be accomplished using lower layer security services.

From a security perspective, only the SCPS-SP implementation need be *trusted* (in a security sense) to always perform its function correctly. The protocol layers above and below SCPS-SP are not relied upon for security services and therefore do not have to be trusted to perform correct security processing. The SCPS-SP is the *funnel* through which all classified or sensitive data must pass to become non-sensitive (via the correct application of confidentiality services). It is also the *funnel* through which the data must pass for its authenticity to be determined.

The SCPS-SP provides the security services confidentiality, integrity, and authentication, as defined in Table 5-1. The specification for SCPS-SP is provided in a CCSDS Recommendation, reference [4].

Table 5-1: Security Services

Service	Definition	Comments
Confidentiality	Protection of data from unintended or inappropriate disclosure	Typically applied to classified or sensitive (e.g., mission critical) data Service is provided via encryption (encryption algorithms are <i>not</i> specified in SCPS-SP) Provides high assurance of protection
Integrity	Protection of data against unauthorized modification or destruction	The receiver is assured to receive <i>exactly</i> what the sender transmitted Provides the ability to detect if tampering has been attempted during transmission
Authentication	Guarantee of the identity of the source of an action.	Two types of authentication: - Source Authentication—assurance of the identity of the source of data - User Authentication—assurance of the identity of an individual user (user authentication is an application layer service and is not part of the SCPS-SP specification)

5.2 SCPS-SP HERITAGE

The SCPS-SP is a bit-efficient hybrid based on several other security protocols such as:

- a) the Secure Data Network System (SDNS) "SP3" protocol,
- b) the ISO Network Layer Security Protocol (NLSP),
- c) the Integrated Network Layer Security Protocol (I-NLSP), and
- d) the Internet Engineering Task Force's (IETF) Internet Protocol Security (IPSEC) Encapsulating Security Payload (ESP) and Authentication Header (AH) protocols.

All of these security protocols provide data confidentiality, data integrity, authorization, and access control, but not bit efficiency. They have large headers and thus far greater overhead than is acceptable over space or other band-width constrained networks. The SCPS-SP has been refined in order to ensure minimal transmitted bit-overhead. However, SCPS-SP does not support the generality of services provided by its less efficient predecessors.

5.3 MISSION APPLICATIONS OF SCPS-SP

Applications of the SCPS-SP to space missions are shown in table 5-2. SCPS-SP does *not* assume the use any particular cryptography, algorithms, or key management. Each mission or system is free to adopt a security policy that provides an appropriate degree of protection for its data. SCPS-SP provides end-to-end confidentiality, integrity, and authentication, but does not protect against traffic analysis or the interception of enciphered data. Those missions requiring such prevention must also use link or physical layer protection services.

Table 5-2: Space Mission Applications of SCPS-SP

Security Service	Application	Examples & Comments
Data Confidentiality	Telecommanding Security	Ensure that commands sent to a spacecraft are protected from unauthorized disclosure and modification
	Payload Data Protection	Ensure that proprietary information is protected from unauthorized disclosure: e.g., imagery to be sold for profit
Data Integrity	Telecommanding Security	Ensure that the command received by the spacecraft is exactly the command that was transmitted from the ground
	Payload Data Protection	Ensure that the data received on the ground is the exact data transmitted by the spacecraft
Authentication	Telecommanding Security	Ensure that only an <i>authorized</i> location (e.g., control center) can command a spacecraft or an instrument
	Payload Data Protection	Ensure that only an <i>authorized</i> location (e.g., control center) can retrieve data transmitted by a spacecraft

5.4 SCPS-SP PROTECTION METHODS

The SCPS-SP encapsulates transport protocol data units (TPDU) into a security protocol data unit (SP-PDU). The TPDU may be enciphered to provide confidentiality, may have an Integrity Check Value (ICV) calculated and appended to provide integrity (non-forgability) of the TPDU, or both. Explicit authentication requires the use of either the integrity and/or the confidentiality services. In the case where both integrity and confidentiality are both required, integrity is applied first and then confidentiality. The specific confidentiality and integrity algorithms are local choices based on the requirements of the local security policy.

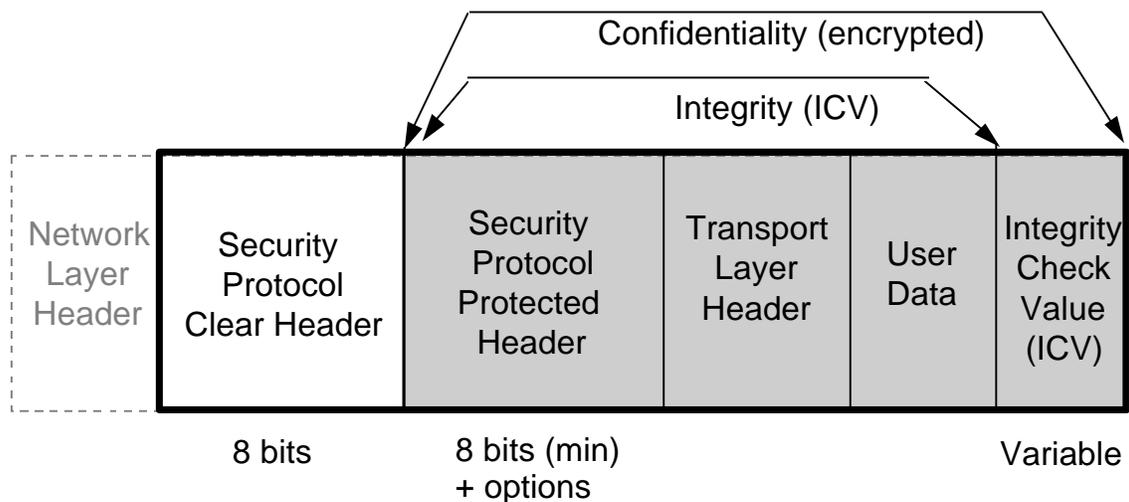


Figure 5-2: Security Header Layout

The SCPS-SP employs both a clear and a protected text header as illustrated in Figure 5-2. The clear header, which must remain un-enciphered, provides routing information to the security protocol. The protected header contains information which may be enciphered along with the user data (e.g., upper layer protocol headers plus user payload data), depending upon the system security policy being enforced by the SCPS-SP as well as the user's security services request. The security protection which the SCPS-SP attempts to provide is derived from a combination of the protection requested by the SCPS-SP user and the protection imposed by the security domain administrator. For example, a user might not request confidentiality services, however the local system security policy might require that all data transmitted employ confidentiality. From a security administration perspective, the enforcement of the local security policy should have precedence over user-based security requests.

Although the degree of protection afforded by some security mechanisms is dependent on the use of some specific cryptographic or secure hash techniques, correct operation of the Security Protocol is not dependent on the choice of any particular encipherment, decipherment, or integrity algorithm. Algorithm choices are left as a local matter as a function of protection requirements and security policy.

In the same vein, neither the choice nor the implementation of a specific security policy is

within the scope of the Security Protocol specification. The choice of a specific security policy, and therefore the protection that will be afforded by the use of SCPS-SP, is left as a local matter.

5.5 SECURITY ASSOCIATION ATTRIBUTES

The SCPS-SP is assumed to have access to a Security Association (SA) database which is a repository of information necessary for the secure operation of the SCPS-SP. The SA database shall be indexed based on a source and destination address pair. The source and destination address pair form a Security Association Identifier (SAID). For bit efficiency, the SAID is not transmitted in the SCPS-SP header. Associated with each entry in the SA database are attributes for the processing of the S-PDU such as the encipher key, the key length, the key expiration, the initialization vector (IV) length, the encipherment algorithm, the integrity algorithm, and the Integrity Check Value (ICV) length. These attributes are used by the authentication and encipherment services to provide end-to-end security for protocol data units.

The security processing attributes in the SA database are either manually pre-placed in the communicating systems employing the Security Protocol or negotiated via a Security Association (SA-P) or Key Management Protocol (KM-P) such as the Internet Engineering Task Force's (IETF) Internet Security Association and Key Management Protocol (ISAKMP). The negotiation protocol is beyond the scope of the Security Protocol specification and is therefore not specified.

5.6 SECURITY PROTOCOL OPERATION

The operation of SCPS-SP is illustrated in Figure 5-3 and described in Sections 5.6.1 and 5.6.2 below.

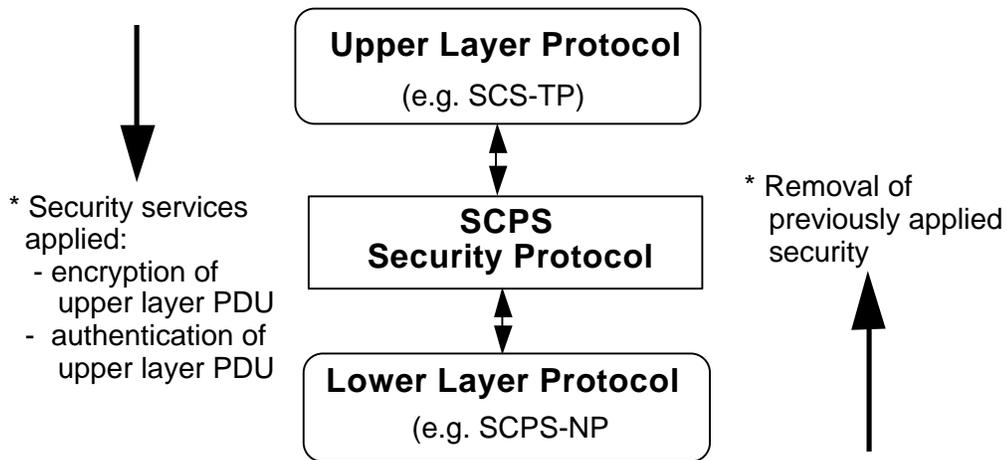


Figure 5-3: SCPS-SP Operation

5.6.1 SCPS-SP TRANSMISSION FUNCTIONS

At the sending end, the Security Protocol:

- receives a PDU from an upper layer protocol (e.g., SCPS-TP),
- attempts to identify a Security Association (SA) database entry based on source and destination addresses. If an existing SA entry is *not* found, an SA (or Key Management) Protocol must first establish an association (if manual, pre-placement of attributes is used, the SA database entry must already exist),
- applies requested (or required, per security policy) security services (e.g., confidentiality, integrity, authentication),
- sends the PDU to the next lower protocol (e.g., SCPS-NP, IP) for transmission over the network

5.6.2 SCPS-SP RECEPTION FUNCTIONS

At the receiving end, the Security Protocol:

- Receives a PDU from a lower layer protocol (e.g. SCPS-NP, IP),
- Identifies SA database entry or, if one is not found, discards PDU,
- Based on security attributes in the SA database, decides what actions to take:
 - Decipher the PDU, and/or,
 - Check the PDU integrity, and/or,
 - Authenticate the explicit source address, and/or,
 - Check the classification of an explicit label against the range of classification allowed on the connection by the SA,
- Pass the PDU to next upper layer protocol (e.g., SCPS-TP).

5.7 END-SYSTEM (ES) TO INTERMEDIATE-SYSTEM (IS) INTERACTIONS

The SCPS-SP can operate between SCPS-SP end systems (ES) on an end-to-end basis. However, “real-life” implementations may use a security front-end device (e.g. an intermediate system (IS)) that implements SCPS-SP and provides security services for an enclave of non-SCPS-SP systems. An example of this might be a cluster of instruments on-board a spacecraft. The spacecraft’s central data handler (CDH) might implement the security protocol to provide end-to-end protection where the spacecraft endpoint is not the instrument but rather the vehicle itself. Figure 5-4 shows an example of such a configuration.

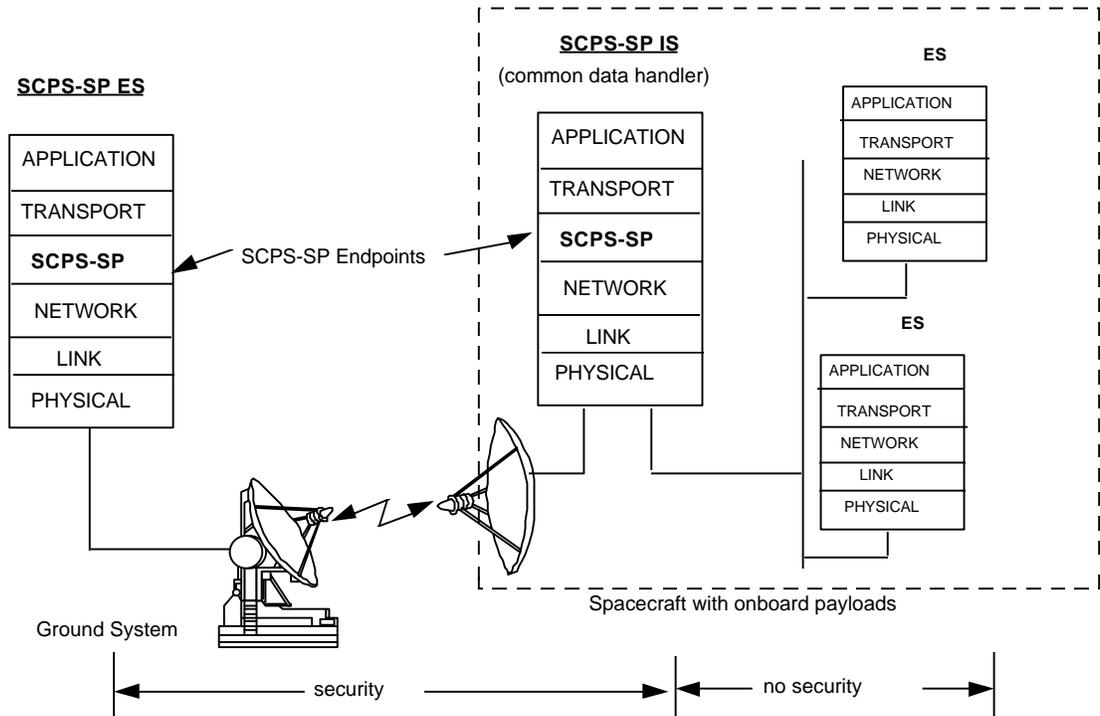


Figure 5-4: SCPS-SP in an On-board Intermediate System

6 SCPS TRANSPORT PROTOCOL (SCPS-TP)

6.1 SCPS-TP OVERVIEW

Transport layer protocols support end-to-end communication between applications using the connectivity provided by an underlying network. The SCPS Transport layer corresponds with Layer 4 of the Basic Reference Model for Open Systems Interconnect (Reference [?]), as shown in figure 6-1. It is designed to meet the needs of a broad range of space missions. The SCPS-TP was developed because existing transport layer protocols do not provide acceptable performance in the space mission environment.

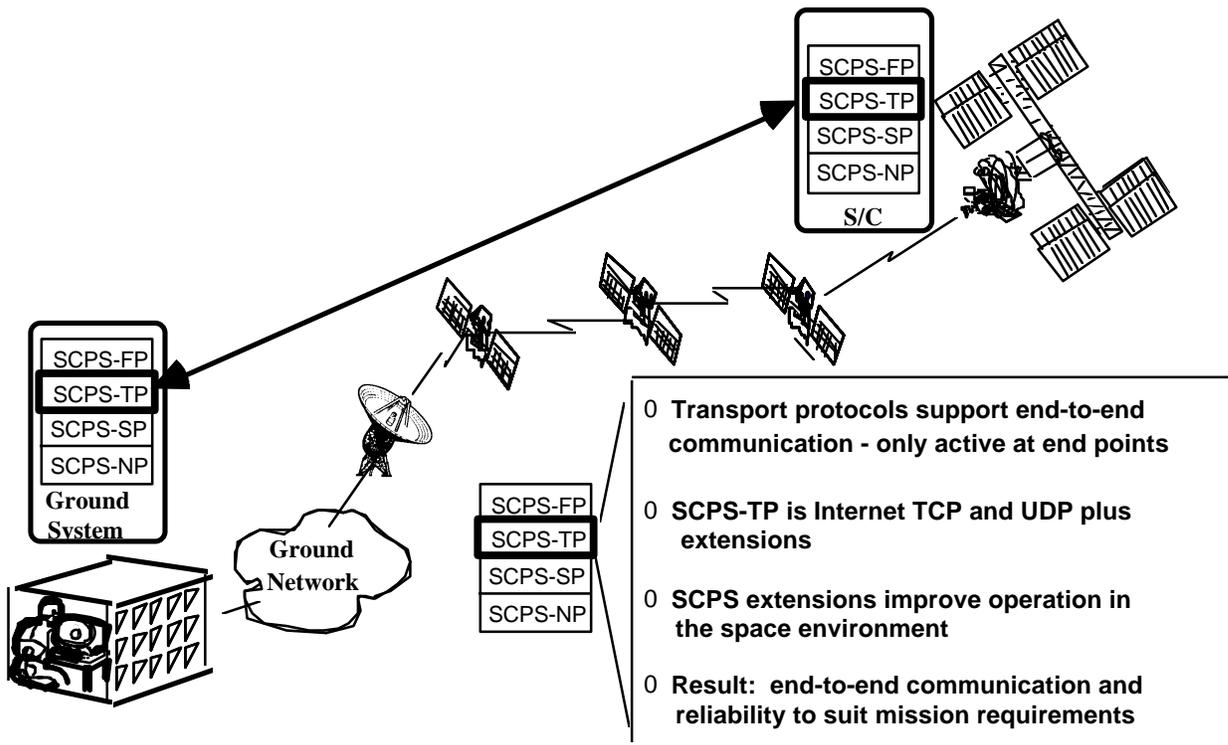


Figure 6-1: The SCPS Transport Protocol

The SCPS-TP is based on the Internet Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) (RFC 793 and RFC 768, respectively), the Internet Host Requirements Document, section 4 (RFC 1122), and the “TCP Extensions for High Performance” (RFC 1323). SCPS extensions to these base standards improve performance in the space environment by addressing the communications difficulties and resource constraints encountered in space missions. Unmodified TCP performs poorly in communications scenarios with large Bandwidth-Delay product (high rate with moderate to high delay, or moderate to high rate with long delay), and cannot operate efficiently with the unbalanced links typical of space-ground communications (See Annex E for comparative test results). The Specification of SCPS-TP is provided in a CCSDS Recommendation, reference [5].

The SCPS Transport Protocol is designed to support operation in current and future space communication environments. The modifications to the base protocols are intended to address the communication environments and resource constraints that systems fielded into these environments typically face.

Several technical requirements were allocated to the transport layer of the SCPS protocol suite:

- a) support for communication with full reliability, partial reliability, and minimal reliability;
- b) efficient operation in a wide range of delay, bandwidth, and error conditions;
- c) efficient operation in space-based processing environments;
- d) support for precedence (priority);
- e) support for connectionless multicasting;
- f) support for packet-oriented applications.

The SCPS Transport Protocol (SCPS-TP) refers collectively to the protocols that provide the full reliability, best-effort reliability, and minimal reliability services. Both the full reliability and partial reliability services are provided by enhancements and minor modifications to Transmission Control Protocol (TCP). The minimal reliability service is provided by the User Datagram Protocol (UDP).

6.2 SCPS-TP REQUIREMENTS

The SCPS-TP is designed to support end-to-end communication and reliability to suit mission requirements within current communication environments and those of upcoming missions. These mission requirements include:

- a) Communication with full reliability, best-effort transfer, or minimal reliability
- b) Efficient operation in a wide range of delay, bandwidth, and error conditions
- c) Efficient operation in space-based processing environments
- d) Support for precedence (priority) based handling
- e) Support for connectionless multicasting
- f) Support for packet-oriented applications

SCPS-TP can support end-to-end communication across connections with end points at a variety of locations to meet the varying needs of specific missions. An endpoint on the spacecraft may be connected to an endpoint in a ground station, which provides an application gateway to the true, remote endpoint, or may be connected (via network layer) directly to the remote endpoint. Detailed functional requirements for SCPS-TP are listed in Appendix C of this Report.

6.2.1 DIFFERENCES BETWEEN COMMUNICATIONS ENVIRONMENTS

The Transmission Control Protocol (TCP) provides an excellent base of technology for extension. It is a highly robust protocol, widely distributed, and is freely available.

Hundreds of individuals, world-wide, work to ensure that TCP continues to meet the needs of the Internet community. The Internet community currently employs a terrestrial communication environment, and TCP is optimized to provide service to this environment. The space and mobile communication environments may have a number of similar communication characteristics to the terrestrial environment by virtue of operating across terrestrial networks as part of the end-to-end network. However, there are significant differences between the terrestrial and space environments that affect communication protocol performance. It should be noted that many of the characteristics of the space environment are also characteristic of mobile and wireless communication. As a result, many of the SCPS enhancements may be applicable to the mobile and wireless communication community.

Table 6-1 presents a summary of the main factors that affect TCP performance when operating in the space or mobile communications environments.

Table 6-1. Factors Affecting TCP Performance in Non-Terrestrial Environments

Factor	Terrestrial Communication	Space and/or Mobile Communication
Bit-Error Rate	Typically $< 10^{-9}$	10^{-4} to 10^{-12}
Round-trip Delay	Milliseconds to seconds	Seconds to hours
Continuity of connectivity	Continuous	Intermittent
Forward and Reverse Link Data Rates	Symmetric	10:1 to 1000:1 forward to reverse link data rate ratio
CPU and Memory Capacity	Relatively large	Relatively small
Communication Goals	Fair access over time High aggregate throughput High reliability	Maximum throughput during contact period Maximum link utilization Selectable reliability level
Primary Sources of Data Loss	Congestion	Congestion Corruption Link Outage

The following paragraphs discuss these factors.

6.2.1.1 Bit-Error Rates

The error performance of typical terrestrial networks has improved to a point that it is no longer considered as a typical source of data loss. With sufficient channel coding and application of radiated power, some tactical and satellite links can approach the error performance of terrestrial networks. However, this is not the typical case, especially in situations in which the power, weight, and volume of the communications gear is constrained.

The loss of data due to bit-errors has a disproportionately bad effect on TCP performance because TCP interprets any loss as an indication of network congestion. The appropriate response to network congestion is to reduce the offered load to the network. TCP's congestion response reduces the offered load by half, then builds back slowly over several subsequent round trips. The effect of this in response to bit-errors is to significantly underutilize the communication channel.

6.2.1.2 Round Trip Delay

Round trip delays in the terrestrial communication environment are typically in the tens of milliseconds to low hundreds of milliseconds. (Round trips across the continental United States average approximately one hundred milliseconds.) In the spacecraft communication environment, round trip times of five hundred milliseconds are the minimum that one expects when communicating through a geostationary satellite, with each hop through a satellite adding another five hundred milliseconds. Deep space communications can increase round trip delays to hours.

Long round-trip delays limit the usefulness and effectiveness of TCP's (or any closed-loop system's) feedback from the remote communication endpoint. This causes problems when the protocol needs to react to changes in the network, but does not receive feedback about those changes until long after the change has occurred.

Note that long delays are not exclusively a result of speed-of-light propagation times. Low data transmission rates add delay to a network, as can half-duplex operation. Finally, queuing in intermediate systems is a source of delay (and the primary source, in the terrestrial communication environment).

6.2.1.3 Continuity of Connectivity

The terrestrial communication environment can be characterized as a network with a very infrequently-changing topology. Orbiting systems have predictable, but possibly highly dynamic, connectivity characteristics. Low Earth Orbiting satellites typically have connectivity through a single ground station 10% of the time or less. Changes to the number of ground stations or the satellite's orbit can improve this, but even NASA's Tracking and Data Relay Satellite System (TDRSS) offers only about 90% coverage. Further, tactical systems have unpredictable connectivity characteristics, due to system mobility and potential system mortality.

6.2.1.4 Forward and Reverse Link Capacity

In the terrestrial communication environment, communication links are typically full duplex with the same data rate in both directions. This is not the case in space environments. Rather, it is not unusual to have large differences in forward and reverse link capacities. Ratios of 1000:1 are not unusual. This degree of asymmetry causes problems for TCP, which uses a stream of acknowledgments as a self-clocking mechanism for transmitting data packets. Thus, very-low-capacity acknowledgment channels limit the transmission rate of data packets.

6.2.1.5 CPU and Memory Capacity

In the terrestrial communications environment, the availability of computing resources is essentially unrestricted. This is not the case in spacecraft where power, weight, and volume are all precious commodities. The amount of computational resource available to any subsystem in a spacecraft must be traded off against the benefits of applying that resource elsewhere. Therefore, it is important to be aware of these constraints. Note that restrictions on power may affect other factors listed here. Notably, power restrictions may increase error rates, decrease data rates (increasing delay), and may affect continuity of connectivity.

6.2.1.6 Communications Goals

A major TCP goal is to provide users fair access to the network over time. By fair access we mean that no single user can monopolize a communication channel when others need to use it. TCP also attempts to provide high aggregate throughput, and provides high reliability.

These communication goals are good. However, the space and mobile communications environment may explicitly NOT wish to provide fair access to the communication resources. Rather, access may need to be on a strict precedence basis, high precedence users being given priority over resources at the expense of lower precedence users.

Further, TCP does not assume that maximization of link utilization is a priority. It intentionally under-drives the link at the beginning of a connection and after loss, in an attempt to determine the sustainable capacity of the link.

Finally, TCP offers a fully-reliable service, preserving completeness, sequence, and correctness. TCP trades delay (incurred as a result of retransmission) and buffer space to provide these features. Its companion protocol, the User Datagram Protocol (UDP), provides an unreliable service, with no preservation of sequence or completeness. However, for some types of data, such as image data, a partial-reliability service that preserves sequence and correctness, but possibly not completeness, may be appropriate. In the case of image data, the idea is that the possible loss of a single scan line (or a part of a scan line) should not significantly delay the delivery of the remainder of the image, but that the order of the scan lines is important to preserve.

6.2.1.7 Primary Source of Data Loss

As previously mentioned, data loss due to bit-errors and to topological instability is rare in the terrestrial environment. The primary source of loss in terrestrial networks is congestion, and TCP is optimized to control congestion. The space and mobile communications environment are mixed-loss environments, with losses occurring due to all three causes: bit-errors, topology changes (link outages), and congestion. To treat all losses as congestion results in unnecessary reductions in offered load. The increased round trip times in these environments delays the restoration of full-rate transmission.

6.3 SCPS-TP SERVICES

SCPS-TP refers collectively to the protocols that provide full reliability, best-effort reliability, and minimal reliability services end to end. The full reliability service is provided by TCP. The best-effort service is provided by TCP with minor modifications. The minimal reliability service is provided by UDP. Table 6-2 describes these services.

Table 6-2: SCPS TP Transport Services

Service	Definition	Comments
Reliable Transfer (TCP)	End to end data transfer of a sequence of data units with full reliability (complete, correct, in sequence, no duplication)	Uses sequence checking to assure sequence and avoid duplication. Uses acknowledgments and retransmission requests to provide completeness. Closes connections without loss of data Completeness is guaranteed, so missing data that cannot be "filled in" by retransmission results in undelivered data beyond the gap.
Best Effort Transfer (BETS)	Transfer with "best effort" reliability (correct, in sequence, no duplication, possibly with gaps). Permits delivery of data with gaps caused by inability to obtain retransmissions.	Under good conditions, BETS provides the same service as TCP, but with the option to continue receiving data if conditions deteriorate. Spacecraft often have limited on-board buffer space to hold data until reception is acknowledged from the ground. Limited contact time can cause gaps (pending retransmissions) at the end of a contact.
Unreliable Transfer (UDP)	Connectionless. Sends data in datagrams. Transfer with minimal reliability (correct, possibly incomplete, possibly out of sequence).	No Sequence numbering; no acknowledgment of receipt; no retransmissions. Similar to current telemetry protocols.

6.4 SCPS-TP EXTENSIONS TO TCP

The SCPS-TP addresses the environmental constraints described above with a number of different extensions and modifications to TCP:

These enhancements are summarized in Table 6-3, and described in the following paragraphs.

Table 6-3 SCPS-TP Modifications to TCP to Address Communication Problems

Factor	Tactical and/or Space Communication	SCPS-TP Modifications
Bit-Error Rate	10^{-4} to 10^{-12}	Corruption response SNACK Header compression
Round-trip Delay	Seconds to hours	Window scaling Timer modifications
Continuity of connectivity	Intermittent	Link outage support
Forward and Reverse Link Data Rates	10:1 to 1000:1 forward to reverse link data rate ratio	Rate control Ack frequency reduction Header compression
CPU Capacity and Memory Availability	Restricted	Header pre computation Record boundaries
Communication Goals	Maximum throughput during contact period Maximum link utilization Selectable reliability	Congestion control optional (rate control to support) Header pre computation Separate corruption response SNACK Partial-reliability operation
Primary Sources of Data Loss	Congestion Corruption Link Outage	Separate response per loss type SCMP signaling Configurable default source of loss

6.4.2 BIT-ERROR RATES

SCPS-TP has developed three capabilities to address the possibility of data loss due to bit-errors. The first is an explicit response to corruption, rather than congestion, as a cause of loss. The second is the Selective Negative Acknowledgment (SNACK) capability. The third is the loss-tolerant header compression mechanisms.

6.4.2.1 Explicit Corruption Response

When TCP responds to an isolated data loss, it reduces its transmission rate by half and doubles its retransmission timer. SCPS-TP's response to corruption does neither of these things. Rather, both the transmission rate (controlled by the congestion window) and the retransmission time out value remain unchanged.

6.4.2.2 Selective Negative Acknowledgment

The SCPS-TP Selective Negative Acknowledgment (SNACK) capability has been developed to identify specific data that requires retransmission, and to request immediate retransmission of that data. The SNACK capability is invoked when the receiver creates and transmits the

SNACK option to the data sender on a regular acknowledgment. The receiver does not send the SNACK option immediately upon detecting a data loss, in case packets have become misordered within the network.

6.4.2.3 SCPS-TP Header Compression

SCPS-TP defines a header compression capability to reduce the size of transmitted packets. This header compression capability operates at the endpoints of the SCPS-TP connection. As a result, headers are only compressed once, regardless of the number of hops that the data requires. Further, this header compression scheme is loss-tolerant, meaning that the loss of one packet does not render subsequent packets unintelligible.

6.4.3 ROUND TRIP DELAY

SCPS-TP addresses the problems imposed by round-trip delay with two capabilities - one defined by the Internet community and one defined in SCPS-TP.

6.4.3.1 Window Scaling

The Window Scaling option (defined in RFC 1323) permits TCP to have more than 64k bytes of data outstanding (unacknowledged) at one time. (Note that at T1 data rates, 1.54 Mbps, a one-half second round trip delay would result in over 96k bytes of data outstanding.) The window scaling option simply imposes a scaling factor to the advertised window, increasing the maximum data that could be outstanding by powers of two, up to 2^{13} .

6.4.3.2 Timer Modifications

SCPS-TP increases the range of typical TCP timers to allow round trip delays of minutes to hours. Further, SCPS-TP initializes its retransmission timer based on data from the routing structure. This allows routes to remote systems to be configured with a reasonable initial estimate of the round-trip time, thus avoiding retransmission time-outs at the beginning of a connection.

6.4.4 CONTINUITY OF CONNECTIVITY

SCPS-TP depends on signaling from the network layer (the SCPS Network Protocol's SCPS Control Message Protocol, SCMP) to identify link outages. This permits SCPS-TP to differentiate between link outages and other causes of packet loss.

6.4.4.1 Signaling of Link Outages

The SCPS Control Message Protocol entity depends on information from local link interfaces (for example, a satellite communications channel) to determine whether the link is available or not. Such information can be inferred from, for example, through scheduling information or from explicit data-link layer signaling. The SCPS Network Protocol entity maintains simple state information about the availability of outbound links. When the link's status changes (for example, from "available" to "unavailable"), SCMP sends a signal indicating the change to recent users of that link. This SCMP signal is received by the SCMP entity at the data source. If, in the case of a link transition from "available" to

“unavailable,” another route to the destination cannot be identified, the “link out” signal is passed up to SCPS-TP.

6.4.4.2 Link Outage Support in SCPS-TP

When the SCPS-TP receives a message from its local SCMP entity that a link is out, it ceases to transmit new data. Additionally, it stops its normal retransmission timers and periodically “probes” the link to determine if it has been restored. These probes are either packets with a single byte of data, or they are acknowledgments (if there is no data waiting to be transmitted). The transmission of a probe is not counted as a retransmission of data, so the connection will not be terminated as a result of exceeding the maximum retransmission count. When SCPS-TP receives an indication that the remote entity is again reachable, either through new packets being received from the remote SCPS-TP or from an SCMP message indicating that the link is restored, it resumes its normal mode of operation.

6.4.5 FORWARD AND REVERSE LINK CAPACITY

Operation of TCP over highly-asymmetric channels tends to result in sustained under utilization of the high-capacity channel, as mentioned above. This is a result of TCP’s use of acknowledgments as clocking mechanisms for transmitting data. SCPS-TP has three capabilities that work together to improve the utilization of the high capacity channel: rate control, acknowledgment frequency reduction, and header compression.

6.4.5.1 Rate Control

SCPS-TP provides a rate control mechanism to “spread” the transmission of data across a time interval, replacing TCP’s acknowledgment-clocking mechanism. SCPS-TP uses a “token-bucket” rate control mechanism, with the rate control parameters associated with a particular route. All SCPS-TP users on a single host that share that route share the capacity of that route. The rate control also provides a means of limiting the rate of transmission of acknowledgments, something that TCP cannot do.

6.4.5.2 Acknowledgment Frequency Reduction

TCP attempts to acknowledge at least every-other packet that is received. If TCP detects that a packet is missing, it sends an acknowledgment for every packet. As previously mentioned, limitations on acknowledgment channel capacity result in under utilization of the data channel. SCPS-TP breaks the dependency on acknowledgments as clocking mechanisms, and therefore allows the acknowledgment rate to be reduced.

SCPS-TP permits the user to explicitly specify the rate at which acknowledgments will be sent. If channel capacity permits, this rate should be at least twice per round trip.

6.4.5.3 SCPS-TP Header Compression

SCPS-TP header compression reduces the size of SCPS-TP headers. By reducing the size of acknowledgments, the load on the (low data rate) acknowledgment is correspondingly reduced. SCPS-TP header compression may be enabled or disabled on a per-route basis.

6.4.6 CPU AND MEMORY CAPACITY

While current TCP implementations tend to be efficient in their use of CPU and memory resources, SCPS-TP has implemented some further enhancements that take advantage of the environment. These enhancements are header pre computation, the provision of record boundaries, and the implementation of some memory-efficient buffering strategies. Only the record boundary modification is strictly a *protocol* feature (meaning that it has end-to-end significance). The other enhancements are local implementation issues, and do not require the cooperation of the remote system.

6.4.6.1 Header Pre Computation

The SCPS-TP implementation provides a header pre computation capability to improve CPU use. Its application is in situations where data collection takes place over long periods of time compared to the time when the communication link is available, and it assumes that the data will be transmitted at high rates once the link becomes available. On an existing connection, when the link becomes unavailable, SCPS-TP continues to accept data from the user (to the limits of its available memory), and does all protocol processing possible. When the link becomes available, the timing-related protocol processing is performed and the queued data is transmitted. The effect of the header pre computation is to amortize the bulk of the protocol processing across the time that the link is unavailable, reducing the “spike” in processing required when the link becomes available.

This capability is completely implementation-dependent. There is no protocol mechanism required to support it. Further, in some situations (for example, high-rate data acquisition and low-rate data transmission) it is inappropriate for use. However, in the case of sustained observation and bursty transmission, header pre computation smoothes the CPU utilization over the two periods. The reference implementation of SCPS-TP provides header pre computation as its intrinsic behavior.

6.4.6.2 Record Boundaries

TCP provides a byte-stream-oriented transmission capability. That is, it does not guarantee the preservation of record boundaries from end to end. This forces applications to provide their own application-layer framing mechanisms to delimit their data units. SCPS-TP provides a record boundary option that does this application data delimiting function. This results in a memory savings when two or more applications have implemented independent application-layer framing software.

6.4.6.3 Memory Buffer Strategies

The reference implementation of SCPS-TP provides memory management that is optimized for efficient use of memory. This is the intrinsic behavior, and requires no user action.

6.4.7 COMMUNICATION GOALS

SCPS-TP addresses the communication goals of the space and tactical communication environments with five enhancements to TCP. To address the goals of maximizing throughput and link utilization during a contact period, TCP’s congestion control

mechanisms are made optional. Header pre computation reduces the protocol processing required at the time that the link is available. And SCPS-TP's corruption response and SNACK capabilities maintain high link-utilization when experiencing bit-errors.

SCPS-TP defines a partial reliability service to address the goal of selectable reliability.

6.4.7.1 Optional Congestion Control

SCPS-TP makes optional the standard congestion control capabilities within TCP. However, if TCP congestion control is not enabled, system designers must ensure that congestion is either controlled by other means or is not possible in the network due to resource reservation.

6.4.7.2 Header Pre Computation

Header pre computation can improve link utilization by reducing the amount of protocol processing required during the time that the link is available. (This benefit accrues if the processor is the performance bottleneck in the system.)

6.4.7.3 Separate Corruption Response

The SCPS-TP corruption response improves link utilization by not interpreting data loss due to bit-errors as data loss due to congestion. When responding to corruption, the transmission rate (and therefore link utilization) is not reduced.

6.4.7.4 Selective Negative Acknowledgment

The Selective Negative Acknowledgment (SNACK) capability improves link utilization by providing a means to unambiguously identify and request immediate retransmission of missing data.

6.4.7.5 Partial-reliability Service

SCPS-TP provides a partial reliability service, called BETS, to ensure correct, in-sequence, but possibly incomplete data delivery. When the BETS capability is enabled, SCPS-TP on the sending side attempts to retransmit packets a user-specified number of times, then continues on as if the packets had been acknowledged (rather than aborting the connection, as standard TCP does). If no retransmissions are desired, the sender discards the packet after its initial transmission. At the receive side, the receiving SCPS-TP entity waits for retransmissions until its receive buffers fill to a user-specified level, then the missing data is signaled to the user. After receiving the signal that a block of data is missing, the receiver can continue reading data beyond that block.

6.4.8 PRIMARY SOURCES OF DATA LOSS

SCPS-TP addresses mixed-loss environments by providing the ability to respond to different types of loss with responses that are appropriate for that type. The SCPS Control Message Protocol (SCMP) provides signaling mechanisms to inform SCPS-TP about the types of loss being experienced. Finally, the SCPS-TP default response can be configured to invoke either the congestion response or the corruption response.

6.4.8.1 Separate Responses for Each Type of Loss

SCPS-TP has separate responses for congestion, corruption, and link outages. The congestion response is the same as that in TCP.

6.4.8.2 SCMP Signaling for Different Loss Types

The SCMP provides separate signals for congestion (the “source quench” signal), corruption (the “corruption experienced” signal), and link outage (the “link out” and “link redirect” signals). Upon receipt of these signals, SCMP informs SCPS-TP and updates local state information.

6.4.8.3 Configurable Default Source of Loss

If SCMP cannot determine the cause of loss or the signal does not reach SCPS-TP, SCPS-TP must invoke some response to that loss. The reference implementation of SCPS-TP allows each route to be configured with a default source of loss, congestion or corruption. If no signals are received indicating the cause of loss, SCPS-TP will invoke its default response.

6.5 SPACE MISSION APPLICATIONS OF SCPS-TP

The SCPS-TP protocols are small in implementation size, with options that can be enabled or disabled by system administration/management action to meet the needs of the specific operating environment. Table 6-4 provides examples of space mission applications of SCPS-TP services.

Table 6-4: Space Mission Applications of SCPS-TP

Transport Service	Applications	Examples & Comments
Reliable Transfer (TCP)	Uplink	Uplink of individual commands, command sequences.
	Downlink	Transfer of any data that is not oversampled, especially compressed data
Best Effort Transfer (BETS)	Downlink	Any payload data that is useful despite occasional gaps. Limited on-board resources and intermittent connectivity make BETS the best choice for most payload data, unless the data is useless unless complete.
Unreliable Transfer (UDP)	Uplink	Contingency operations. Commanding in the blind.
	Downlink	Repetitive or oversampled data. Used when minimal delay is more important than completeness.

6.6 SCPS-TP SUMMARY

SCPS-TP is based on TCP and UDP and can be interoperable with commercial TCP and UDP protocol products. SCPS-TP provides extensions to TCP to address the space mission communication environment. These optional features (extensions) may be enabled or disabled to meet the requirements of specific missions.

7 SCPS FILE PROTOCOL (SCPS-FP)

7.1 OVERVIEW OF SCPS-FP

The SCPS-FP is located in the application layer, as shown in Figure 7-1, and uses the transport services of SCPS-TP or internet TCP. The Specification of SCPS-FP is provided in a CCSDS Recommendation, Reference [6].

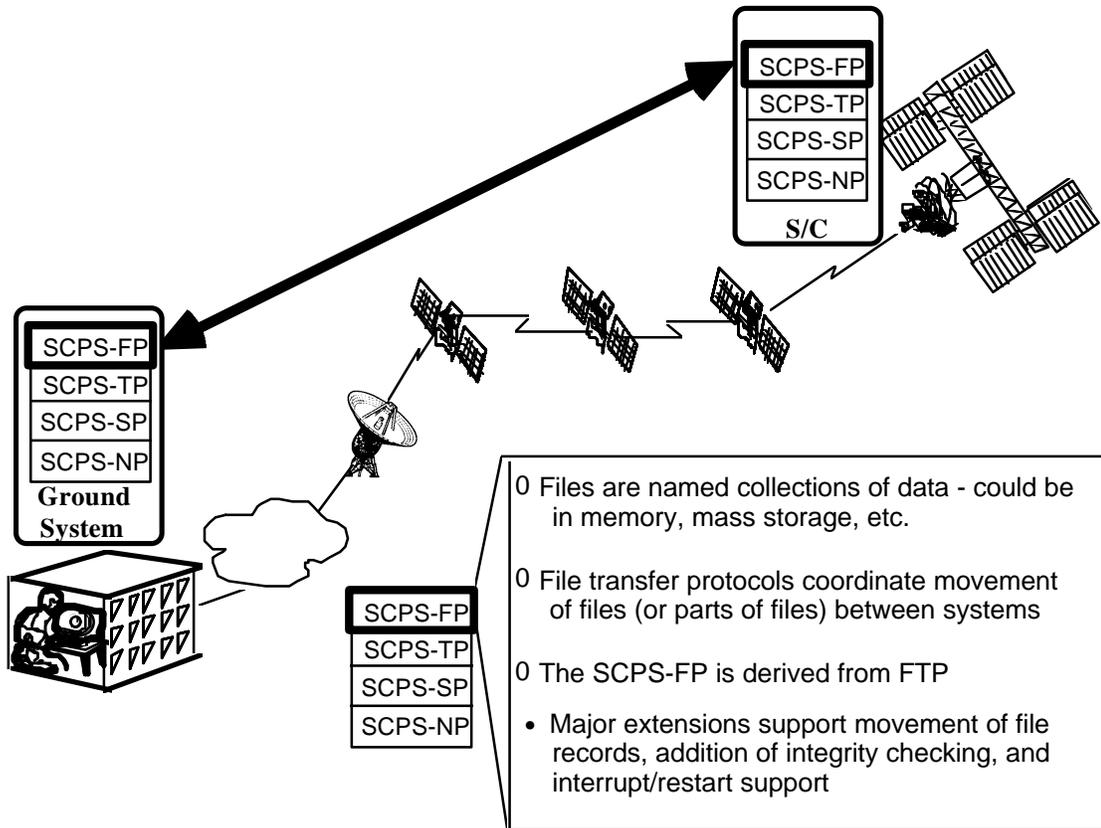


Figure 7-1: SCPS File Transfer Protocol

The SCPS-FP is designed to support the file transfer and file operation requirements of current and future space missions. The SCPS-FP is derived from the Internet File Transfer Protocol (FTP), and uses the FTP model (Figure 7-2) as described in RFC 959. Like FTP, SCPS-FP operates between end systems that may use different file storage and access techniques. SCPS-FP implementations can be tailored to meet the requirements of a range of missions, and is interoperable with FTP.

7.2 SCPS-FP SERVICES

7.2.1 SCPS-FP HERITAGE

The objectives of the Internet FTP protocol [Reference 18] , from which SCPS-FP is derived, are:

- a) Promote sharing of files (computer programs and/or data)
- b) Encourage indirect or implicit (via programs) use of remote computers
- c) Shield a user from variations in file storage systems among hosts
- d) Transfer data reliably and efficiently

The model for FTP, as defined in RFC-959, Section 2.3 [18], is shown in Figure 7-2. In the figure, the protocol interpreter is indicated by PI, and the data transfer process is indicated by DTP. FTP uses two transport connections, a control connection to exchange control information, and a data transfer connection to move file data.

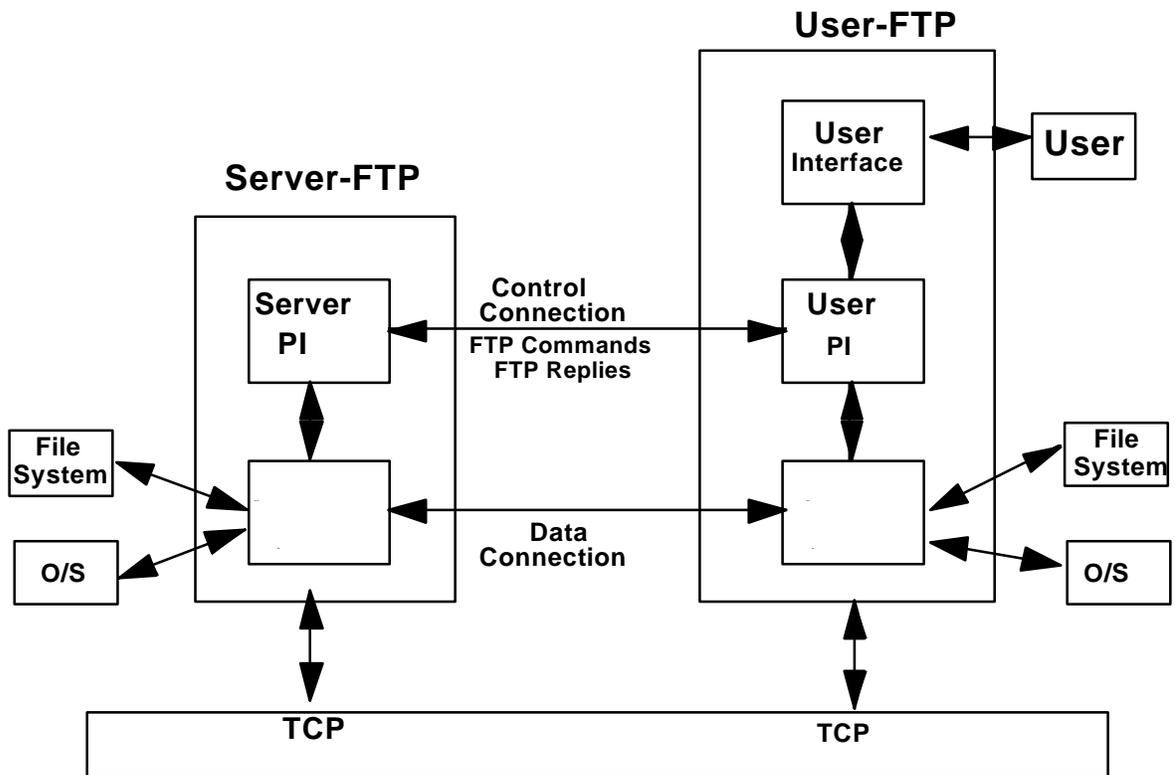


Figure 7-2: Internet FTP Model

7.2.2 SCPS-FP DATA REPRESENTATION AND STORAGE

Data is transferred from a storage device in the sending host to a storage device in the receiving host. Often it is necessary to perform transformations on the data because data storage representations in the two systems are different.

In addition to different representation types, FP allows the structure of a file to be specified. Three file structures are defined in FP:

- a) file-structure, where there is no internal structure and the file is considered to be a continuous sequence of data bytes,
- b) record-structure, where the file is made up of sequential records, and
- c) page-structure, where the file is made up of independent indexed pages.

7.3 SCPS-FP REQUIREMENTS

File services required by space missions include:

- a) transfers of command and data files to spacecraft,
- b) transfers of application software to spacecraft,
- c) transfers of science or mission data to ground without special processing to reorder or merge data sets,
- d) limited management of files onboard spacecraft (delete, rename, and directory services),
- e) automatic restart of transfers after an interruption,
- f) read portions of files resident onboard spacecraft, and
- g) make updates/changes to files onboard spacecraft, without sending a complete replacement for the file to make minor modifications.

The detailed functional requirements for SCPS-FP are listed in Appendix C of this report.

7.4 SCPS-FP MODIFICATIONS TO FTP

The base standards, RFCs 959 and 1123, provide for basic command/reply dialog and much of the functionality required for the SCPS-FP. However, some additional functions are needed for space mission operations, and some standard FTP functions are unnecessary and inefficient in this environment. Thus SCPS-FP provides the following modifications to FTP:

- a) Enhanced error recovery and restart capabilities
- b) Manual interrupt
- c) Record read
- d) Record update
- e) Enhanced file integrity features
- f) Suppression of reply text

Each of these modifications is described below.

7.4.1 ENHANCED ERROR RECOVERY AND RESTART

Interruption of service is much more common in space environments than in terrestrial environments due to the fact that the hosts are in motion relative to each other. SCPS-FP provides a restart procedure to make it easier to restart an interrupted transfer.

7.4.1.1 Automatic Restart

Automatic restart provides the capability to restart failed file transfers with no user intervention.

7.4.1.2 Manual Interrupt And Manual Restart

Manual interrupt provides the capability for the User-FP to temporarily stop a file transfer and then restart it at a later time.

7.4.2 RECORD READ

Record Read provides the capability to read part of a file resident on a remote system rather than read the complete file. The status and contents of the file at the server site are unaffected by the record read service.

7.4.3 RECORD UPDATE

The Record Update extension provides the capability to update or change part of a file on board a spacecraft without transferring the entire file. The record update data consists of an update script that indicates which records to delete and modify in the remote file and where to add new records. A checksum computed against a local copy of the remote file is transmitted along with the remote file names, and the update script to the Server-FP as control data.

At the receiving end, the Server-FP verifies that the remote target file is the correct file to modify by comparing the remote file's checksum with that provided in the control data, applies the update script to the remote file, and stores the result at the Server-FP in a new file. If the Server-FP is unable to perform the operation (e.g., because of an invalid script), the requester is notified.

7.4.4 ENHANCED FILE INTEGRITY

For the purposes of the SCPS-FP, there are two aspects of integrity:

- a) Ensuring that the data sent is the data received, which is provided by the reliable transport service of SCPS-TP, and
- b) Ensuring that when a transfer (or record operation) fails, any intermediate changes are undone (or rolled back). This aspect of integrity is not provided by FTP and thus has been added to SCPS-FP.

7.4.4.1 File Transfer Integrity

The User-FP and Server-FP rollback any incomplete changes made to a file during an operation that has terminated with errors, thereby restoring the affected file to its pre-transfer state. The method and sophistication of the rollback mechanism is left as a local implementation issue.

7.4.4.2 Record Data Integrity

The User-FP and Server-FP rollback any incomplete changes made to a file during a read or update operation that has terminated with errors, thereby restoring the affected file to its pre-record operation state. The method and sophistication of the rollback mechanism is left as a local implementation issue.

The Record Update operation is also are required to verify the contents of the destination file using a CRC to ensure that a user does not update the wrong file inadvertently. The User-FP and Server-FP employ a CRC checksum for the record update service to ensure the data integrity of the accessed files.

7.4.5 SUPPRESSION OF REPLY TEXT

The Suppress Reply Text option provides the capability to suppress the reply text from being sent by the server except for replies that must be parsed by the client. The server responds to the client (user) only with the reply code in the syntax specified in RFC 959. The SCPS-FP server avoids the overhead of text replies and maintains interoperability with COTS FTP.

7.5 CONFORMING IMPLEMENTATIONS OF THE SCPS-FP

There are four types of implementations of the SCPS-FP:

- a) **Minimum FP**—a conforming minimum SCPS-FP implementation, which provides basic file transfer capabilities (e.g., for a severely resource constrained mission).
- b) **Full FP**—a conforming full SCPS-FP implementation, which provides the file services needed by most space missions.
- c) **Full FP + minimum FTP**—a conforming full SCPS-FP implementation (as in ‘b’) above, plus services required for minimal compatibility with Internet's FTP (i.e., support of the minimum command set specified in RFC-959).
- d) **Full FP + minimum FTP + optional FTP**—implementations that address other, optional, non-SCPS features of FTP, in addition to the services described in ‘c’) above.

The details of the protocol requirements and the service interface for each type of conforming implementation are provided in the SCPS-FP Specification [Ref. 6].

The SCPS-FP assumes that a reliable transport service (the SCPS-TP or Internet TCP) is available for connection establishment and management and for data transmission.

ANNEX A

GLOSSARY

Connection: A connection is defined by information that is named, persistent, and shared across the systems supporting an instance of communication. For transport protocols, these systems are the endpoints that terminate the transport protocol, but not intermediate systems.

End System: An addressable network entity within the SCPS Network.

Gateway: A network-addressable system that terminates a protocol at a given layer and invokes similar services at the same layer of an adjacent network..

Host: A network-addressable system that may send or receive network-layer packets, but does not forward packets.

Maximum Segment Size: The maximum amount of user data that can be carried in a Segment. This value is calculated by subtracting the size of the network, security, and transport layer headers from the MTU size.

Maximum Transmission Unit: The Maximum Transmission Unit (MTU) specifies the maximum amount of data that the subnetwork layer will accept in a single subnetwork service request. The MTU for a route is the minimum of all known MTUs along that route. (Note: It is anticipated that this value will be known and managed as part of the routing table information, however, techniques for dynamically discovering the MTU of a route exist. Refer to RFC 1191, "Path MTU Discovery" for more information.)

Router: A network-addressable system that may send, receive, or forward network-layer packets.

Segment: The Protocol Data Unit of the Transmission Control Protocol (TCP).

Service-Access-Point: A point at which the services of a layer are made available to the layer above it.

ANNEX B

SCPS Green Book ACRONYMS

This Annex provides an identification of the acronyms used in this SCPS Green Book.

AOS	Advanced Orbiting Systems
CDH	Central Data Handling
DOD	Department of Defense
FP	File Protocol
FTAM	ISO file transfer and access mechanism
FTP	File Transfer Protocol
ICMP	Internet Control Message Protocol
ICV	Integrity Check Value
IP	Internet Protocol
IS	End-System (ES) to Intermediate-System
IV	Initialization Vector
MIB	Management Information Base
NLSP	Internet Network layer security protocols
NP	Network Protocol
OSI	Open Systems Interconnection
PDU	Protocol Data Unit
PVC	Permanent Virtual Circuit
RFC	Request for Comments
SA	Security Association
SCMP	SCPS Control Message Protocol
SCPS	Space Communications Protocol Standards
SDNS	Secure Data Network
SEU	Single Event Upsets
SP	Security Protocol
TCM	Terminology, Conventions, and Methodology
TCP	Internet Transmission Control Protocol
TP	Transport Protocol
TPDU	Transport Protocol Data Units
UDP	User Datagram Protocol

ANNEX C

FREQUENTLY ASKED QUESTIONS

This Section answers questions that have often been asked during SCPS briefings and design reviews, and in CCSDS Red Book reviews of the SCPS specifications. It covers the rationale for development of the SCPS, the requirements and constraints that guided that development, and the applicability of the protocols to future space missions. Most of this is covered in the body of the CCSDS Report: *SCPS Concept, Rationale, and Application Notes*, but this question and answer format provides a less formal discussion of these topics, and serves as a starting point for an on-line FAQ file as additional questions arise and experience with the protocols accumulates.

C1 PURPOSE AND SCOPE

C1.1 Q: Why were the SCPS protocols developed?

The principal goal of the SCPS effort was to lower lifecycle costs by reducing development and operations costs in space communications systems.

The SCPS program was initiated in response to several demands:

- a) A need for standard protocols to support reliable data transfer.
- b) The need to accommodate evolving multi-node mission configurations that require in-space network routing.
- c) The need to significantly reduce operations costs and thus maintain the ability to produce results from space missions in the face of decreasing funding.
- d) The need to provide compatibility and interoperability with the internet.

The SCPS are designed to meet these demands by increasing standardization and interoperability both within and among CCSDS Agencies and other developers and operators of spacecraft.

C1.2 Q: Do SCPS replace earlier CCSDS protocols?

No. SCPS augments these protocols by providing reliable stream or file transfer over CCSDS frames at the link layer and dynamic networking for those missions that need it. CCSDS Packets remain as the telemetry and telecommand source message format. CCSDS RF and modulation Recommendations are also unaffected.

NOTE – At the time of publication, efforts were underway to enhance CCSDS link protocols to support multiplexing of SCPS Network Protocol and other network-level data units into CCSDS virtual channel frames.

Some missions may choose not to use some features of earlier CCSDS protocols, if their needs are better met by SCPS. For example, a mission that uses SCPS network addressing over a CCSDS frames may find it unnecessary to use several virtual channels for data routing.

C1.3 Q: What types of missions are the SCPS intended to support?

SCPS is aimed at a broad range of space missions including:

- a) Support for spacecraft in low-earth and geosynchronous orbits, as well as lunar and planetary spacecraft. The primary emphasis has been on support of missions at lunar distances or closer. SCPS network and security protocols are relatively immune to communications delay, and thus can support deep-space missions today. SCPS Phase 3, beginning in 1997, will provide additional capabilities for data transport and file transfer in deep-space missions.
- b) Support of spacecraft with a range of on-board communication and on-board data handling resources, including those with limited on-board computer and memory resources, as well as those with multiple, high-capacity on-board computers with extensive data storage.
- c) Support of multi-node in-space networks, including:
 - 1) LEO spacecraft constellations.
 - 2) Cluster-like missions.
 - 3) Deep space orbiter/lander/rover planetary missions.

C2. APPLICABILITY

C2.1 Q: Does a mission have to use all of the SCPS protocols?

No. Each mission can choose the appropriate layers and the appropriate options within those layers. The individual protocols provide flexibility and optional features that allow designers to tailor a communications protocol to meet the requirements and constraints of a mission, without extensive software development.

C2.2 Q: Doesn't the availability of so many options defeat the purpose of a standard?

Not really. Like most protocol stacks, SCPS provides three classes of options. First, spacecraft designers can choose which of the protocols to use. A mission might be designed to operate without one of the SCPS protocols, for various reasons. The second class of options is the choice of features to include in a specific copy of a protocol entity, to tailor that copy for use in a particular environment. Such adaptation is often made through compile-time options. Finally, some options do not affect size or overall capability of a protocol implementation, but are simply setup parameters that configure the run-time protocol entity to optimize performance or provide compatibility with peer protocol entities.

Each of the options was designed to accommodate the real and unavoidable differences from one mission to another—differences in objectives, hardware, environment, or operations. These differences must be accommodated in some way. Use of a few carefully designed options within a family of protocols will lead to lower cost and lower risk than the multiple, uncoordinated, point solutions that would otherwise be necessary.

C3. RELATIONSHIP TO OTHER PROTOCOLS

C3.1 Q: Can SCPS be used with the CCSDS link layer protocols? Specifically, packet telemetry downlink and CCSDS telecommand uplink?

Yes. The SCPS were designed to operate over CCSDS space-ground links, although use of other link layer protocols is possible. The figure below shows just one example of a protocol profile in which SCPS operates over CCSDS telemetry and telecommand frames. In this example, on-board applications can both send and receive CCSDS Packets. This data can be transferred as streams of packets (via SCPS-TP) or as files of packets (via SCPS-FP). End-to-end data protection is provided by SCPS-SP, and network addressing is provided by SCPS-NP. The use of an expected enhancement to telemetry multiplexing is assumed; this enhancement would provide for multiplexing of SCPS-NP or other network PDUs (e.g., internet IP packets) into the data zone of CCSDS telemetry frames (VCDUs).

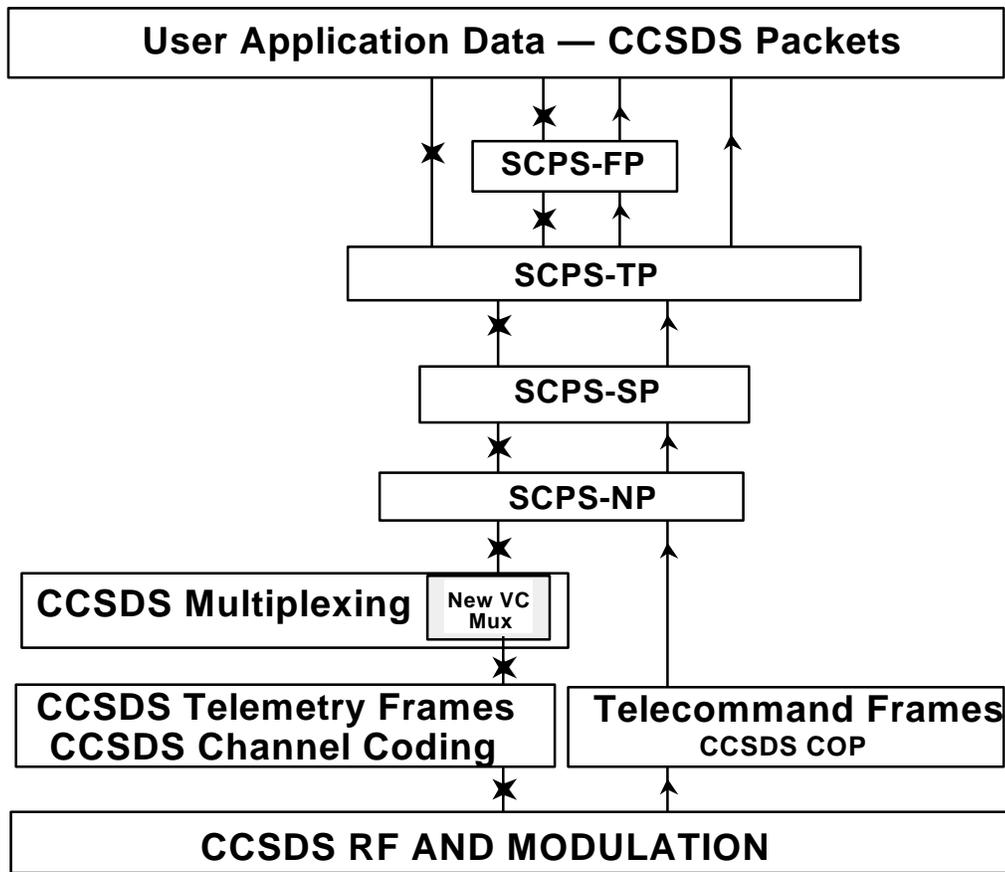


Figure C3.1: One Example of SCPS supported by CCSDS Link Protocols

C3.2 Q: How are the SCPS and CCSDS protocols related to internet and OSI protocols?

In the figure below, OSI and the internet protocols—FTP, TCP, and IP—are shown along with CCSDS and SCPS. Within the CCSDS/SCPS family this figure shows some profile options not shown in the previous figure. For example, the option of running SCPS over the AOS path (packet) service via a convergence layer is shown.

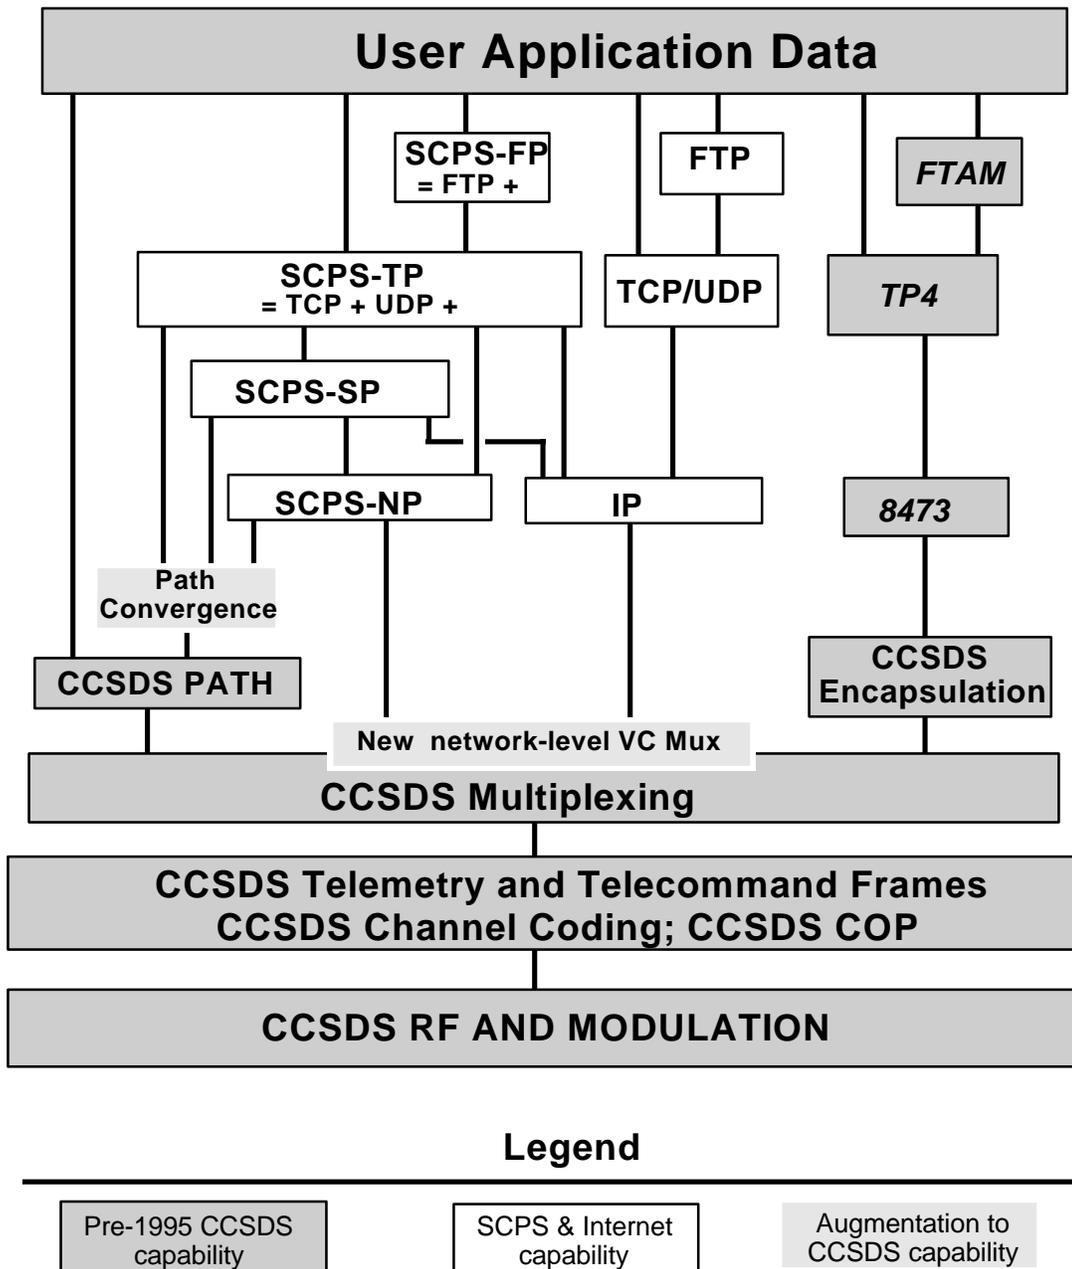


Figure C3.2: SCPS, TCP/IP, and CCSDS Protocol Profiles

C3.3 Q: Which typical mission data operations tasks are performed by the SCPS and CCSDS space link protocols ?

The table below shows a number of tasks performed in most space missions, and lists the protocols that may be used to support the task.

Task	CCSDS	SCPS
Real-time data stream	PT-SP, AOS-Path	SCPS-TP (UDP option)
Real-time in-sequence data stream	PT-SP, AOS-Path	SCPS-TP (best-effort option; no re-xmit)
Near-real-time, nearly complete data	<i>M*</i>	SCPS-TP (best-effort option)
reliable transfer of telemetry	<i>M*</i>	SCPS-TP (complete option)
guaranteed delivery of event-driven messages	<i>M*</i>	SCPS-TP
uploading software files to subsystems/payloads	<i>M*</i>	SCPS-FP
uploading software patches to subsystems/payloads (record operations)	<i>M*</i>	SCPS-FP
uploading data files and memory contents to subsystems and payloads	<i>M*</i>	SCPS-FP
downloading data files and memory contents from subsystems and payloads	<i>M*</i>	SCPS-FP
transmission of commands, or sequences of commands, from the ground user to the subsystems/payloads	Telecommand	SCPS-TP
initiation, control, and monitoring of stored command programs	Telecommand	SCPS-TP
interactive access from an onboard terminal to a ground data system or from a ground terminal to an onboard data system		SCPS-TP
data base access and file transfer from ground to crew workstations		SCPS-FP
secure end-to-end data transfer		SCPS-SP

**M => Accomplished through manual operations and/or mission-specific software.*

C3.4 Q: Won't the Reed-Solomon and other FEC capabilities of the CCSDS link provide nearly perfect delivery without the overhead of transport and file transfer protocols?

There are several parts to this answer.

- a) Not all space links use such powerful error correction.
- b) Even with Reed-Solomon error correction, performance is better described as “usually perfect, but with occasional gaps”—almost all data is delivered, and the delivered data is error-free, but occasionally whole messages are lost because a link layer frame is not decodable, due to weather effects or operational problems.
- c) Many user data traverse multiple sub-networks—not just the space link. Data can be lost during transit over LANs in end systems or in gateways due to buffer overflow. Trying to solve a high-level problem—Did all of it get to the final destination correctly?—at lower levels is futile². No amount of confirmation, correction, or protection on a link-by-link or subnet-by-subnet level will do the job. On the other hand, it is not possible to communicate effectively without some reliability measures at lower layers, given the link characteristics and intermittent connectivity in space operations. A balance of upper-layer, confirmed, end-to-end services supported by links good enough to avoid excessive retransmission is the optimal solution.

C4 SECURITY

C4.1 Q: What is the scope of the protection provided by the SCPS Security Protocol?

The SCPS Security Protocol (SCPS-SP) provides *end-to-end* security services and resides between OSI layers three (network) and four (transport). SCPS-SP provides confidentiality (encryption), integrity, and authentication services. Access control is provided as a by-product of confidentiality and authentication. Protocol data units (PDUs) can use confidentiality and integrity independent of one another, or combined. Authentication must be used with either integrity, confidentiality, or both. Confidentiality ensures that the data is protected from eavesdropping from its source to its destination. Integrity prevents unauthorized modification of the data while it is in transit from its source to its destination. Authentication provides assurance to the receiver that the data actually came from the claimed source.

C4.2 Q: Why can't we just use link encryption?

Link encryption is a powerful security mechanism, but can only be applied on a hop-by-hop basis. If there are multiple communication hops involved between a source and a destination, then link encryption devices are required between each hop and the data must be decrypted when received and then re-encrypted for transmission onward, potentially exposing the data to those who should not have access to it, or to undetectable modification or corruption.

For example, if data is to be sent from an instrument control facility, to a spacecraft control facility, to a ground station, and across a space link to a spacecraft, then link encryption devices would be required on each end of the communications circuit between the instrument

²Saltzer, J., et al. *End-To-End Arguments in System Design*, 1984 [24]

control facility and the spacecraft control facility. Additional link encryption devices would be required on the communications link between the spacecraft control facility and the ground station. Finally, additional link encryption devices would be required on the space link - in the ground station and aboard the spacecraft. However, the data would be decrypted (and then re-encrypted) at the spacecraft control facility and at the ground station. Therefore, the data is potentially available to those who should not see it, and also exposed to potential corruption.

Link encryption provides total obscurity of the information on the link whereas all of the header information in the layers below SCPS-SP is openly available when using SCPS-SP. If prevention of traffic analysis or data interception is required, only link encryption can provide these protection services. A combination of SCPS-SP (to provide end-to-end protection) and link encryption (on selected vulnerable links) provides the most secure solution.

C4.3 Q: Can't these security protocols lock me out of my own spacecraft?

No, not if the spacecraft was designed for use with the Security Protocol. All spacecraft potentially have the problem that on-board resources (e.g., an on-board computer) may not operate correctly resulting in control and operations problems. In order to ensure emergency control of the spacecraft, the hardware command decoder should execute several critical functions directly in its own hardware (e.g., critical actuators, on-board computer re-initialization) rather than relying on other *upstream* spacecraft resources. In this manner, emergency commanding will be performed by a hardware function directly behind the spacecraft's radio receiver and before any other on-board computer resources. However, from a security perspective, this results in a potential vulnerability to the spacecraft unless the emergency commands are protected in some manner. A tradeoff must be made by the designers between spacecraft emergency safety and overall spacecraft security.

C5 SCPS DESIGN CHOICES

C5.1 Q: Why were the SCPS protocols developed by modifying internet protocols rather than by revising CCSDS protocols?

A primary goal of the SCPS effort was to extend internet connectivity into space. The rationale for this approach is that both the data systems and the personnel (designers, operators, users) associated with space missions are already using internet protocols. The communications services that they need in space are very similar to those they have in ground networks. The easiest, lowest risk, and most direct way to achieve this goal was to adapt the protocols that are used on the ground.

Previous CCSDS protocols were not designed to provide the functionality that the SCPS offer. CCSDS protocols used for return (or downlink) data provide error-protected, sequenced data streams. This service supports real-time data acquisition and quick look analysis. It also makes possible the production of best-effort (nearly complete) data sets from multiple dumps of data. But these protocols were not intended to support automatic, real-time retransmission to provide complete or best-effort data streams, or to provide reliable file transfer. Adding these services would require additional protocol layers and

complexity equal to the SCPS approach, but would not yield the benefit of internet compatibility, nor capitalize on the vast experience with internet protocol development and use.

C5.2 Q: Why were modifications to internet protocols needed?

Although the internet protocols provide an excellent basis for space communications protocol development, the space environment presents a number of constraints that are seldom encountered in the design of terrestrial data communications networks. These include:

Physical differences, including:

- a) Space link delays ranging from milliseconds to hours.
- b) Potentially noisy space data links.
- c) Limited space link bandwidth.
- d) Variation in sub-network types from simple busses to local and wide area networks.
- e) Interruptions in the end-to-end data path that can vary from single bits lost due to high background noise or single event upsets (SEUs), momentary link interruptions caused by intense bursts of noise, and longer interruptions caused by spacecraft antenna obscurations.

Operational differences, including:

- a) The inherently sporadic nature of contact between space and ground.
- b) "Teleoperations" activities may pose a maximum latency requirement.

Resource differences, including:

- a) Limited onboard processing power.
- b) Limited onboard program memory.
- c) Limited onboard data buffering.
- d) Extreme asymmetry in bandwidth between forward and return links.

Except for a very narrow range of operational conditions, the current off-the-shelf, internet protocols do not satisfy the requirements encountered in the space mission environment. SCPS adopted a policy of using of COTS-supported standards wherever possible, to capitalize on established user interface familiarity and minimize software development costs. This approach also mitigates risk by exploiting the hundreds of thousands of hours of operational experience that the internet protocols have accrued.

C6 PERFORMANCE, OPERATIONAL, AND COST ISSUES

C6.1 Q: In trying to cover the spectrum of space missions, doesn't the SCPS approach sacrifice efficiency?

There is always some penalty in resources or performance for a general solution as compared to one optimized to the needs of a particular project. The general solution, however, has advantages of a wider market to share development, testing, and maintenance costs. Wider use and testing means increased reliability and reduced training and operations costs. Finally, custom solutions often do not have the flexibility needed for contingency operations when missions don't go as planned, or evolve into new missions.

C6.2 Q: What features does SCPS provide to help meet the goal of reducing operations costs?

- a) Operate end-to-end.
- b) Provide a service that is consistent and reliable across all networks and missions.
- c) Facilitate automation of space operations.
- d) Reduce integration and test effort.
- d) Provides familiar interfaces and operational paradigms.

C6.3 Q: Have the SCPS Protocols ever been tested over real space links?

Yes, they have. A satellite relay (bent-pipe) experiment to test SCPS-TP was carried out using a US Department of Defense satellite. SCPS- TP performed well, maintaining between 82% and 97% of maximum throughput (depending on packet size) at bit-error rates of up to 10^{-5} . As part of this test, the performance of SCPS- TP was compared to that of Transmission Control Protocol (TCP) using a similar configuration in the laboratory. SCPS-TP performance was equivalent to that of TCP at low bit error rates, and significantly better than TCP's at bit-error rates of 10^{-7} or greater.

The UK Defense Research Agency's Space Technology Research Vehicle (STRV) was utilized to exercise the SCPS protocols onboard an orbiting spacecraft. Several of the SCPS protocols (FP, TP, SP) were uploaded to the STRV flight computer and were tested between space and ground under actual flight conditions. Files were uploaded and downloaded between the ground and the STRV via the use of the SCPS File Protocol (SCPS-FP) and the SCPS Transport Protocol (SCPS-TP). SCPS-TP's ability to hold connections across short contact times, cope with high bit error rates on the space communications link, and its ability to provide high throughput were tested. The SCPS Security Protocol (SCPS-SP) was tested in conjunction with SCPS-TP and demonstrated that the SCPS-TP tests could be carried out in a secure environment."

ANNEX D

PROTOCOL FUNCTIONAL REQUIREMENTS [PFR]

D1 SCPS-NP FUNCTIONAL REQUIREMENTS

Req Ref	Requirement Summary	SCPS-NP Ref
N.1	Support for multicasting	
N.1.1	Shall be able to recognize the group destination specified by the user application, provided that such destination is a valid one	3.8.2
N.1.2	Shall be able to select the group address that correctly corresponds to the destination referred to in N.1.1	3.8.1.3
N.1.3	Shall be able to assign proper group addresses to each outgoing packet that requires one	3.4.1 3.4.2
N.1.4	Shall be able to recognize valid group addresses and properly interpret them. Proper interpretation is defined as accurately determining how to route/relay the packets containing such group addresses.	3.9.2 3.9.3.1.5
N.2	Support for multiple routing options	
N.2.1	Shall be able to request the address of its neighboring node(s) from a routing module(s)	3.9.2 3.9.3
N.2.2	Shall be able to select the proper neighboring node for a packet and transmit the packet to that node.	3.9.2 3.9.3
N.2.3	Shall be able to route a packet to a unicast destination.	3.9.2 3.9.3.1.1
N.2.4	Shall be able to route a packet to a multicast destination consisting of one or more end systems.	3.9.2 3.9.3.1.5
N.2.5	Shall be able to flood route a packet to all space-based end systems.	3.9.2 3.9.3.1.3
N.2.6	Shall ensure that a flood routed packet that has been forwarded by a node is not subsequently forwarded by that same node	3.9.2.3 3.9.3.1.3

DRAFT REPORT: SCPS RATIONALE, REQUIREMENTS, AND APPLICATION NOTES

N.3	Packet lifetime support	
N.3.1	Shall be able to assign a maximum-age indication (e.g., hop count or time value) to each outgoing packet that requires one	3.9.1.1 3.9.1.2.4 3.9.1.2.7 3.9.1.2.8
N.3.2	Shall be able to determine the age of an incoming packet and properly interpret it. Proper interpretation is defined as accurately determining whether or not the incoming packet should be discarded due to having reached (or exceeded) its allowed lifetime.	3.9.2.3 (4)
N.3.3	Shall be able to automatically discard a packet which lifetime has been reached (or exceeded)	3.9.2.3 (4)
N.3.4	Shall, when a hop count is in use, be able to properly increment the age of each outgoing packet that requires it (adjusting or recomputing any network layer checksum or forward error correction as necessary).	3.9.2.3 (4)
N.4	Separate reporting of congestion and corruption	
N.4.1	Shall be able to detect and differentiate between network congestion and network data corruption.	3.7
N.4.2	Shall be able to report each of these two conditions to the transport protocol in a way that differentiates between them	3.10.3.2 3.10.3.6
N.4.3	Shall be able to manage and possibly discard data in response to congestion.	3.9.3.2.2
N.4.4	In the event that it is necessary to discard data, data shall be discarded in order from lowest precedence to highest precedence.	3.9.3.2.2
N.5	Support for precedence handling	
N.5.1	Shall be able to recognize the precedence level specified by the application	3.8.2 3.8.3.1.1
N.5.2	Shall be able to provide a default precedence level for those packets that require one	3.8.2
N.5.3	Shall be able to assign the proper precedence level to each outgoing packet that requires one	3.9.1.2 3.9.1.2.9
N.5.4	Shall be able to recognize the precedence level associated with an incoming packet.	3.9.1.2 3.9.1.2.9
N.5.5	Shall be able to process incoming packets in accordance with their assigned precedence level.	3.9.3.2
N.5.6	Shall provide the ability for system configuration personnel to set the default precedence level for a system.	3.8.2 4.1.2
N.5.7	Shall provide for sixteen levels of precedence	3.8.2

DRAFT REPORT: SCPS RATIONALE, REQUIREMENTS, AND APPLICATION NOTES

N.6	Differentiation between real and exercise data	
N.6.1	Shall be able to recognize the data type (real vs. non real) specified by the application	3.8.2 (Precedence)
N.6.2	Shall be able to provide a default data type to each outgoing packet.	3.8.2 (Precedence)
N.6.3	Shall be able to assign the proper data type to each outgoing packet	3.8.2 (Precedence)
N.6.4	Shall be able to recognize the data type associated with an incoming packet.	3.8.2 (Precedence)
N.6.5	Shall be able to process incoming packets in accordance with their assigned data type.	3.8.2 (Precedence)
N.6.6	Shall provide the ability for system configuration personnel to set the default data type for a system	3.8.2 (Precedence)

D2 SCPS-SP FUNCTIONAL REQUIREMENTS

The functional requirements for the protection of data between end points within a space data communications system are stated below.

Requirement P.1: Access control

The SCPS data protection protocol shall provide the capability to control access to network resources. Only those users (or processes acting on behalf of users) with authorization shall be granted access to network resources. Examples of network resources are end systems, transport protocols within an end system, and router.

Requirement P.2: Source authentication

The SCPS data protection protocol shall provide the capability to verify the identity of the end system that originated network communications.

Requirement P.3: Command authentication

The SCPS data protection protocol shall provide a capability to digitally *sign* a message to indicate that the message was actually sent by the user (or process acting on behalf of the user) claiming to send it.

Requirement P.4: Integrity

The SCPS data protection protocol shall provide the capability to ensure that the data sent is exactly the data received. It will provide the assurance that any unauthorized modification of the data will be detected while the data is in transit across the network.

Requirement P.5: Confidentiality

The SCPS data protection protocol shall provide the capability to ensure that the data transmitted across the network can be properly interpreted only by authorized users (or processes acting on behalf of users).

D3 SCPS-TP FUNCTIONAL REQUIREMENTS

The following lists the protocol functional requirements allocated to the SCPS Transport Protocol and the section in the SCPS-TP Specification that addresses each requirement.

Req Ref	Requirement Summary	SCPS-TP Ref
T.1	Full Reliability. Qualifier "Provided that there is end-to-end link availability and sufficient link capacity for retransmissions, the SCPS transport protocol"	
T.1.1	Shall provide the capability to deliver <i>all</i> data segments to the correct destination(s), as addressed at the source.	4.2.1 a 4.2.1 b 4.2.1 c
T.1.2	Shall provide the capability to deliver <i>all</i> data segments in the same order as originated at the source, with no duplicate or extraneous data	4.2.1 a
T.1.3	Shall provide the capability to deliver <i>all</i> data segments for which there are <i>no</i> detected errors	4.2.1 a
T.1.4	Shall provide the capability to recover from detected data transmission errors.	4.2.1 d 4.2.1 f 4.2.1 g
(T.1.5)	Shall support unicast operation (only)	4.2.4.1.4 c
T.2	Best Effort. "Qualifier: Provided that there is end-to-end link availability, the SCPS transport protocol"	
T.2.1	Shall provide the capability to deliver data segments to the correct destination(s), as addressed at the source	4.2.1 b
T.2.2	Shall provide the capability to continue to deliver data segments to the correct destination(s), irrespective of the loss of a subset of the data segments.	5.2.2
T.2.3	Shall provide the capability to deliver data segments in the same order as originated at the source, with <i>no</i> duplicate or extraneous data	4.2.1 a 5.2.2
T.2.4	Shall provide the capability to deliver data segments for which there are <i>no</i> detected errors	5.2.2
(T.2.5)	Shall support unicast operation (only)	4.2.4.1.4 c
T.3	Minimal Reliability	
T.3.1	Shall provide the capability to deliver transmitted data segments to the correct destination(s), as addressed at the source, with no guarantee of (a) order, (b) completeness, or (c) elimination of duplicates	4.3.1
T.3.2	Shall provide the capability to deliver data segments for which there are <i>no</i> detected errors	4.3.2.2 6.1
(T.3.3)	Shall support both unicast and multicast operation	4.3.2.2.3

DRAFT REPORT: SCPS RATIONALE, REQUIREMENTS, AND APPLICATION NOTES

T.4	Multicasting. Shall provide the capability to deliver data segments to any subset of all possible destinations, as addressed at the source, under the minimal reliability transmission criteria.	4.3.2.2.3
T.5	Precedence handling.	
T.5.1	Shall provide the capability to recognize the precedence level specified by the user for a connection ([A] in full reliability and [B] best effort reliability operation) or for a [C] data segment (in minimal reliability operation).	A and B: 4.2.1 h 4.2.4.1.2 b 5.1.2 C: 6.2
T.5.2	Shall provide a default precedence level that can be set by system configuration personnel	7.2.2.4 7.3.2.1
T.5.3	Shall provide the capability to deliver data segments in accordance with their assigned precedence level	4.2.1 h 6.2
T.6	Segmentation	
T.6.1	Shall provide the capability for specification of the maximum segment size, by the system administrator, in accordance with system performance characteristics.	7.1.1
T.6.2	Shall provide the capability for peer transport entities to negotiate a maximum segment size	4.2.1 c
T.6.3	Shall provide the capability to reassemble the finite-sized data segments back into their original form as a unitary message.	4.2.1 c
(T.6.4)	T.6.2 and T.6.3 only apply when employing full reliability and best-effort reliability	
T.7	Operation over wide range of conditions	
T.7.1	Shall be able to be configured to operate in processing environments typical of those available on space-based platforms	8 - Profiles Section To Be Included
T.7.2	Shall be able to support workloads typical of those anticipated for space-based platforms	Impl. rqt (not functional rqt)
T.7.3	Shall be able to operate reliably under the delay, bandwidth, and error conditions typical of space-based communication environments.	4.2.2 4.2.4.2.8 5.4 5.5 5.6
T.8	Graceful closing of connections	
T.8.1	Shall provide the capability to recognize requests for termination of a logical connection originating from the user of that connection	4.2.1 k
T.8.2	Shall provide the capability to recognize termination requests of a logical connection originating from its peer transport entity.	4.2.1 k

DRAFT REPORT: SCPS RATIONALE, REQUIREMENTS, AND APPLICATION NOTES

T.8.3	Shall provide the capability for peer transport entities to mutually agree upon the closure of a logical connection	4.2.1 k
T.8.4	Shall provide the capability to ensure successful delivery of any data segments in transit to a destination prior to the mutually-agreed termination of any logical connection required for that data segment, subject to the caveats expressed in T.1.	4.2.1 k
T.9	Response to congestion and corruption	
T.9.1	Shall provide the capability to differentiate between network congestion and network data corruption, as identified by the network level protocol.	5.4 7.1.9 7.1.10
T.9.2	Shall provide the capability to counteract the identified network congestion anomalies.	5.4
T.9.3	Shall provide the capability to compensate for the identified network data corruption anomalies.	5.4

D4 FILE TRANSFER FUNCTIONAL REQUIREMENTS

The functional requirements for the transfer of data files between end points within a space data communications system are stated below.

Req Ref	Requirement Summary	SCPS-FP Ref
F.1	Operations on entire files	
F.1.1	Shall provide the capability to rename files.	5.1, D7.1 RFC 959 (4.1.3, 5.3.1, 5.3.2, 5.4)
F.1.2	Shall provide the capability delete files.	5.1, D7.2 RFC 959 (4.1.3, 5.3.1, 5.3.2, 5.4)
	If a file directory structure is present in the file system, then the SCPS file transfer protocol (F.1.3 through F.1.6):	
F.1.3	Shall provide the capability to create a directory.	5.1, D7.4 RFC 959 (4.1.3, 5.3.1, 5.3.2, 5.4, App II)
F.1.4	Shall provide the capability to delete a directory.	5.1, D7.5 RFC 959 (4.1.3, 5.3.1, 5.3.2, 5.4, App II)
F.1.5	Shall provide the capability to change the current working directory.	5.1, D7.6 RFC 959 (4.1.3, 5.3.1, 5.3.2, 5.4, App II)
F.1.6	Shall provide the capability to list the names of files in a directory.	5.1, D7.3 RFC 959 (4.1.3, 5.3.1, 5.3.2, 5.4, App II) RFC 1123 4.1.2.7
F.1.7	Shall provide the capability to request the size of a file.	4.1.3, 5.3.1, 5.3.2, 5.4, 5.1, D7.7
F.2	Operations on file records	
F.2.1	Shall provide the capability to read and extract any record or set of records within a file.	5.1, 4.1.3, 5.3.1, 5.3.2, 5.4, D5.1

Req Ref	Requirement Summary	SCPS-FP Ref
F.2.2	Shall provide the capability to insert a record or set of records into any location within a file, where location means at the beginning of a file, at the end of a file, or between other records of a file.	5.1, 4.1.3, 5.3.1, 5.3.2, 5.4, D5.2
F.2.3	Shall provide the capability to replace (overwrite) any record or set of records within a file.	5.1, 4.1.3, 5.3.1, 5.3.2, 5.4, D5.2
F.2.4	Shall provide the capability to delete any record or set of records within a file.	5.1, 4.1.3, 5.3.1, 5.3.2, 5.4, D5.2
	Two party file transfer	
F.3	The SCPS file transfer protocol shall provide the capability for either of two end systems to send and receive a complete file of data.	5.1, D4.1, D4.2 RFC 959 (4.1.3, 5.3.1, 5.3.2, 5.4)
	Proxy file transfer	
F.4	The SCPS file transfer protocol shall provide the capability for either of two end systems to send and receive a complete file of data under the control of a third end system.	5.1, D4.3 RFC 959 (4.1.2, 5.3.1, 5.3.2, 5.4)
F.5	User Initiated Interrupt & Abort	
F.5.1	(Manual Interrupt) Shall provide the capability for the user to cause an interrupt of the data transmission after the start but before the completion of the transfer.	5.1, 4.1.3, 5.3.1, 5.3.2, 5.4, D6.1
F.5.2	(Manual Abort) Shall provide the ability for a user to terminate a file transfer after the start but before the completion of the transfer. An aborted file transfer cannot be resumed.	5.1, D6.2 RFC 959 (4.1.3, 5.3.1, 5.3.2, 5.4)
F.6	System-detected Interrupt Notification	
F.6.1	Shall recognize a notification which identifies that the communications supporting a file transfer has been interrupted. This notification is sent by a lower layer (e.g., the transport layer).	3.5.3
F.6.2	Shall act upon this notification (see F.7.2).	3.5.3
F.7	Resumption After Interrupt	
F.7.1	(Manual Resume): Shall provide the capability to manually resume a file transfer from the point of interruption for manual interrupts and automatically detected interrupts.	3.5, 5.1, 4.1.3, D6.3 RFC 959 (5.3.1, 5.3.2, 5.4)

Req Ref	Requirement Summary	SCPS-FP Ref
F.7.2	(Automatic Resume): Shall provide the capability to automatically resume a file transfer from the point of interruption for automatically detected interrupts.	3.5, 5.1, D2.8, D2.9, D2.10 RFC 959 (4.1.2, 4.1.3, 5.3.1, 5.3.2, 5.4)
	Integrity Over Operations of Entire File	
F.8	The SCPS file transfer protocol shall have the capability to preserve the integrity of operations on entire files. Integrity of operations is defined to mean that the operation performed is the same as the operation requested, and that an operation is not performed upon detection of an error by the file transfer protocol.	3.6.1
	Integrity over operations on file records	
F.9	The SCPS file transfer protocol shall have the capability to preserve the integrity of operations on file records. Integrity of operations is defined to mean that the operation performed is the same as the operation requested, and that an operation is not performed upon detection of an error by the file transfer protocol.	3.6.2
F.10	File Transfer security	
F.10.1	(User Access): Shall provide the capability to restrict user access to the functions of the (file transfer) protocol.	5.1, D3.2, D3.3, D3.4 RFC 959 (4.1.1, 5.3.1, 5.3.2, 5.4)
F.10.2	(File Access): Shall provide the capability to prevent unauthorized access to files.	assume file system to provide capability
	Reply Text Suppression	
F.11	Shall provide the capability to suppress the text from the remote system's reply before the reply is sent.	4.1.3, 5.1, 5.3.1, 5.3.2, 5.4
F.12	Configuration Services	
F.12.1	Shall provide the capability to set the data type for the transfer.	5.1, D2.3 RFC 959 (3.1.1, 4.1.2, 5.3.1, 5.3.2, 5.4)

Req Ref	Requirement Summary	SCPS-FP Ref
	Configuration Services (Continued)	
F.12.2	Shall provide the capability to set the data structure for the transfer.	5.1, D2.4 RFC 959 (3.1.2, 4.1.2, 5.3.1, 5.3.2, 5.4)
F.12.3	Shall provide the capability to set the transmission mode.	5.1, D2.5 RFC 959 (3.4, 4.1.2, 5.3.1, 5.3.2, 5.4)
F.12.4	Shall provide the capability to enable and disable the automatic restart capability.	3.5.2, 4.1.2, 5.1, 5.3.1, 5.3.2, 5.4, D2.10, D2.11
F.12.5	Shall provide the capability to set the maximum number of restarts in an automatic restart.	D2.12
F.12.6	Shall provide the capability to enable and disable the automatic use of the PORT command on transfers.	D2.6, D2.7 RFC 959 (4.1.2, 5.1, 5.3.1, 5.3.2, 5.4)
F.12.7	Shall provide the capability to enable and disable the suppression of the remote system's reply text.	4.1.3, 5.1, 5.3.1, 5.3.2, 5.4, D2.8, D2.9
F.12.8	Shall provide the capability to configure the FP idle timeout for the remote system.	D2.14 RFC 959 (4.1.3, 5.1, 5.3.1, 5.3.2, 5.4)
F.12.9	Shall provide the capability to request the configuration status of the local FP application.	D2.13
F.12.10	Shall provide the capability to request the configuration status of the remote FP application.	D2.13 RFC 959 (4.1.3, 5.1, 5.3.1, 5.3.2, 5.4)
F.12.11	Shall provide the capability to request command help for the local FP application.	D2.15
F.12.12	Shall provide the capability to request command help for the remote FP application.	D2.15 RFC 959 (4.1.3, 5.1, 5.3.1, 5.3.2, 5.4)

Req Ref	Requirement Summary	SCPS-FP Ref
F.13	FP Session Establishment/Termination/Maintenance	
F.13.1	Shall provide the capability to establish and maintain an FP session with the remote system.	D3.1, D3.2 RFC 959 (3.2, 3.3) RFC 1123 (4.1.2.5)
F.13.2	Shall provide the capability to terminate an FP session with the remote system.	D3.4, D3.5 RFC 959 (3.2, 3.3) RFC 1123 (4.1.2.5)
F.13.3	Shall provide notification of the success or failure of a file transfer capability.	4.2, 5.4 RFC 959 (4.2, 5.4) RFC 1123 (4.1.2.11)
F.14	Miscellaneous Services	
F.14.1	Shall provide the capability to execute commands specific to the remote system.	4.1.3, 5.1, 5.3.1, 5.3.2, 5.4, D.8.1 RFC 1123 (4.1.2.8)

ANNEX E

**SPACE COMMUNICATIONS PROTOCOL STANDARDS (SCPS)
BENT-PIPE EXPERIMENT REPORT**

EXECUTIVE SUMMARY

NOTE – This Annex contains the executive summary of a report that documents a test of SCPS over a US military spacecraft. Although the summary contains some background discussion and programmatic status that is no longer relevant, it is included in its entirety to give a snapshot view of the test program.

E1 PURPOSE OF THIS DOCUMENT

This report documents the findings of the SCPS Bent-Pipe test, an experiment using a DOD satellite communications package, designated M-22, to evaluate a set of communications protocols being developed by the joint NASA/DOD Space Communications Protocol Standards (SCPS) project.

E2 BACKGROUND

In the fall of 1992, NASA and the DOD jointly established a technical team (the SCPS Technical Working Group, or "SCPS-TWG") to explore possibilities for developing common space data communications standards. By the end of 1993 the team concluded that wide segments of the US. civil and military space communities have common needs for protocols to support in-flight monitoring and control of civil and military spacecraft.

The program of work to develop these protocols includes specification, simulation, implementation, and testing. The SCPS Bent-Pipe Experiment is the third in a series of tests of the protocols. The first test used simulations to evaluate the protocol functional design. The second test used a laboratory test bed to evaluate the performance and functioning of prototype software that implements the protocols. The third test, the SCPS Bent-Pipe Experiment, used the M-22 system to connect two SCPS end systems together in order to evaluate the performance of one of the protocols over a real space link. Subsequent tests will host the prototype software on different spacecraft in various configurations to evaluate performance and functionality in the intended implementation environment.

After analyzing the protocol functions and other factors, the SCPS-TWG decided that the best use of satellite resources would be to focus the bent-pipe test on the transport layer. The SCPS Transport Protocol (SCPS-TP) is a derivative of Transmission Control Protocol (TCP) with modifications and enhancements intended to improve its performance in a satellite communications environment.

The SCPS Bent-Pipe Experiment was conducted by United States Space Command (USSPACECOM) Directorate of Logistics, USSPACECOM/J4 and MITRE personnel under the direction of, and with support from, the Air Force Materiel Command (AFMC) Space & Missile Systems Center Test and Evaluation office (SMC/TEO), and their contractor personnel.

E3 M-22 TEST OBJECTIVES

The bent-pipe experiment put forth three objectives:

1. To measure the operation of SCPS-TP in a real satellite communications environment.
2. To characterize the performance of SCPS-TP.
3. To compare the performance of SCPS-TP to the Transmission Control Protocol (TCP), the standard Internet reliable transport protocol, in a satellite communications environment.

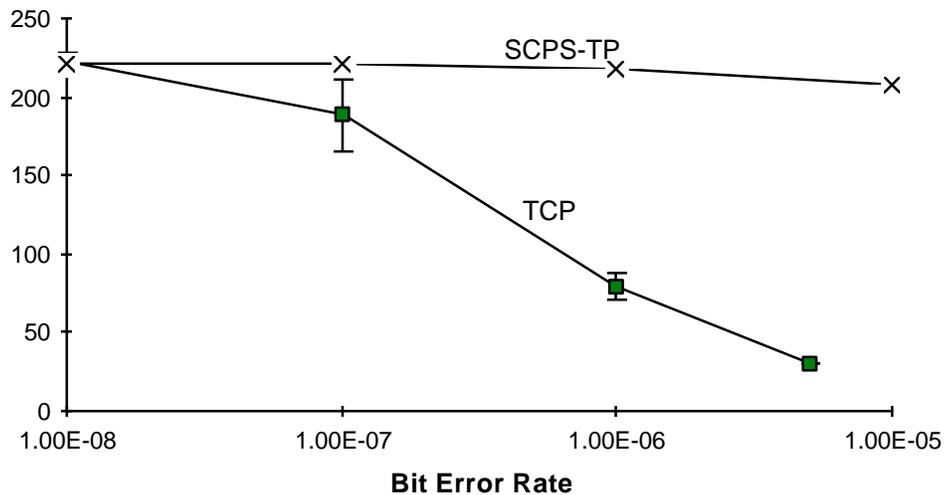
Objectives 1 and 2 were met. Objective 3 was only partially met, since the tests in support of objective 3 could only be conducted in the laboratory environment, due to restrictions of some of

the interface equipment used for the tests.

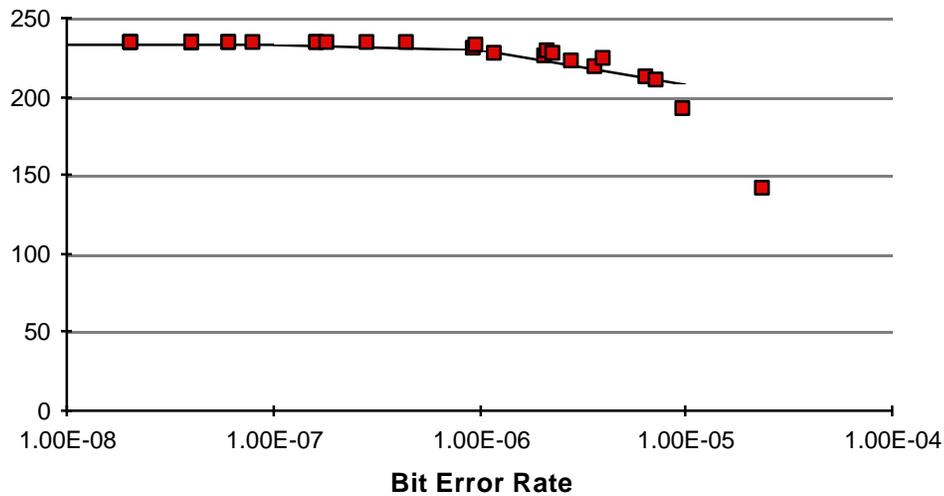
E4 SUMMARY OF RESULTS

The experiment was conducted in two phases: we tested the protocols in the field, and then we repeated the tests in the laboratory. As a result of equipment limitations, we were unable to directly compare the performance of the SCPS-TP with TCP over the satellite link. Rather, we configured a laboratory experiment and conducted the tests there. For TCP comparison testing, we collected throughput data. For the SCPS-TP characterization data, we collected the following measurements: user data throughput, data channel utilization, volume of acknowledgment traffic generated, and bit efficiency.

The following graph compares the throughput of SCPS-TP to that of TCP for a range of bit error rates. This data was collected in the laboratory environment, with the user data per packet being 512 bytes for TCP and 500 bytes for SCPS-TP. The graph shows 90% confidence intervals around each data point. (The confidence intervals for SCPS-TP are closely-spaced, and indistinguishable from the markers for each primary data point.)



The next graph compares the SCPS-TP laboratory data to the bent-pipe experiment data. This graph shows throughput performance for 1000-byte packets. The solid line shows the laboratory data, the individual data points show the M-22 experiment data. (The laboratory data is the mean resulting from averaging five independent tests at each of the following bit-error rates: 10^{-8} , 10^{-7} , 10^{-6} , and 10^{-5} .) It should be noted that not all of the M-22 experiment data corresponds so well to the laboratory data. During the course of the experiment, we identified a number of implementation problems that affected the performance of the protocols. We have corrected most of these, but some are still under investigation. The ability to find and fix these problems was a major benefit of the bent-pipe experiment.



E5 CONCLUSIONS

The bent-pipe experiment provided a beneficial “first step” in SCPS protocol testing. The test provided valuable insight into the performance of SCPS-TP in a real satellite communication environment. In addition, it provided the SCPS developers with an introduction to the type of ground station equipment that may be encountered during operational use of the SCPS protocols. We are able to draw the following conclusions from this activity:

- a) SCPS-TP performed equivalently to TCP in the laboratory simulation of the M-22 environment in the absence of bit errors.
- b) With error rates of 10^{-7} or greater, SCPS-TP showed significantly better throughput performance than TCP (<200 Kbps vs. 30 Kbps at 5×10^{-6}). Limitations of the particular TCP implementation we used prevented testing in the laboratory with a broader range of packet sizes. Limitations of the test hardware in Sunnyvale prevented any testing of TCP over the actual satellite link.
- c) SCPS-TP performs well over a broad range of error conditions:
 - 1) SCPS-TP throughput at a bit-error rate of 10^{-5} was measured at between 82% and 97% of the throughput in the lowest error cases. The variation was a result of different packet sizes.
 - 2) The data channel utilization of SCPS-TP was high when the rate control parameters were properly set. In the laboratory tests, data channel utilization varied between 89% and 96%, while in the bent-pipe test it varied between 70% and 96%. In analyzing the channel utilization information, we identified problems with the rate control strategy. Some of these have been repaired, while others remain under investigation.
 - 3) The amount of acknowledgment traffic generated was low compared to the amount of traffic the acknowledgment channel could accommodate. We improved the acknowledgment strategy to further reduce the amount of acknowledgment traffic generated; the laboratory data reflects those improvements.

- 4) The bit-efficiency of SCPS-TP over IP for 1000-byte and 500-byte packets met the bit efficiency goal of 10% or less protocol overhead at 0 BER, which was established in the requirements phase of the activity. Further, this goal was met for 1000-byte and 500-byte packets with error rates up to approximately 5×10^{-6} .
- b) The testing improved the quality of the SCPS-TP implementation. We made improvements to the acknowledgment strategy and to the rate control implementation as a direct result of the bent-pipe tests. We have identified areas to concentrate upon to ease deployment and field configuration of the protocol.

E6 RECOMMENDATIONS

We offer the following recommendations as a result of the bent-pipe tests.

E6.1 BROAD RECOMMENDATIONS

- a) Continue the program of SCPS testing.
 - 1) Some of the important capabilities within SCPS-TP have not yet been tested outside the laboratory - these include the Best Effort Transport Service, the User Datagram Protocol, and link outage handling. SCPS-TP should be tested in a mixed-loss environment in which losses may be due to congestion, corruption, or link outage.
 - 2) At the time we conducted these tests, we did not have the protocols fully instrumented. While we were able to collect data about the overall characteristics of each protocol test, such as throughput and amount of data transferred, we were not able to gather detailed, instantaneous data. Subsequent protocol tests should be augmented with the ability to gather information about the instantaneous dynamics of protocol operation.
- b) Expand the scope of the testing to more faithfully represent operational communication requirements. The bent-pipe experiment verified the performance of SCPS-TP operating in full reliability mode, in a two-system configuration, with no other traffic present. As testing proceeds, we should attempt to identify a limited set of operational scenarios that represent more realistic uses of the SCPS protocols. Once these scenarios have been identified, we should define and execute a program of experimentation to test the SCPS protocols with these scenarios using as realistic a system configuration as possible.

E6.2 SPECIFIC RECOMMENDATIONS

- a) We recommend that further comparative testing be conducted with an implementation of TCP that permits specification of the maximum segment size. We also recommend that a test environment be identified in which TCP can be compared directly to SCPS-TP over actual satellite communication conditions and that comparative testing be conducted.
- b) We do not currently have instrumentation and analysis tools to allow us to examine the instantaneous data rate on the acknowledgment channel. This is an important capability which we should develop, in order to ensure proper performance over links with restricted acknowledgment channels, and we recommend that it be developed.

ANNEX F

SCPS FP/STRV 1B FLIGHT EXPERIMENT SUMMARY

EXECUTIVE SUMMARY

NOTE – This Annex contains the executive summary of a report that documents a test of SCPS over a U.K. Defence Research Agency spacecraft. Although the summary contains some background discussion and programmatic status that is no longer relevant, it is included in its entirety to give a snapshot view of the test program.

F1 PURPOSE OF THIS DOCUMENT

The Space Communications Protocol Standards - Transport Protocol (SCPS-TP) is being developed by the joint NASA/DOD Space Communications Protocol project. This report documents the findings of the SCPS/STRV Flight Experiment (SSFE) SCPS-TP test.

F2 BACKGROUND

In the fall of 1992, NASA and the DOD jointly established a technical team (the SCPS Technical Working Group, or “SCPS-TWG”) to explore possibilities for developing common space data communications standards. By the end of 1993 the team concluded that wide segments of the US. civil and military space communities have common needs for protocols to support in-flight monitoring and control of civil and military spacecraft. In 1994, the U.K. Defence Research Agency joined the SCPS-TWG with specific interoperability interests for the U.K. Skynet series of military communications satellites.

The program of work to develop these protocols includes specification, simulation, implementation, and testing. The SCPS/STRV Flight Experiment is the latest in a series of tests, that has included simulation, laboratory testing, and a bent-pipe test over a satellite link. The SCPS/STRV Flight Experiment was the first test to actually host the prototype software on a spacecraft, and was intended to evaluate performance and functionality in the anticipated implementation and operational environments.

The protocols tested in the SSFE include the SCPS File Protocol, the SCPS Transport Protocol, and the SCPS Security Protocol. All of the SCPS File Protocol testing made use of the SCPS Transport Protocol, and the SCPS Security Protocol testing used the SCPS Transport Protocol as its data source. The tests of the SCPS File Protocol and SCPS Security Protocol are documented separately (reference [14], [15]).

The SSFE was conducted between 2 January 1996 and 30 April 1996 and between 16 July 1996 and 31 July 1996. The SCPS-TP tests were conducted by U.K. Defence Research Agency personnel stationed at Lasham, England and at Farnborough, England, and by MITRE and Gemini Industries personnel at Reston, Virginia. The tests were conducted at Lasham, England and Reston, Virginia.

F3 SSFE SCPS-TP TEST OBJECTIVES

The objectives of the transport protocol portion of the SSFE were as follows:

- to gain experience in hosting SCPS-TP on an actual spacecraft; and
- to examine the performance of SCPS-TP when running over a real space/ground data link.

In examining the performance of SCPS-TP, we tested three specific capabilities. The following list cites the primary benefits expected from each of the capabilities:

- TCP Timestamps: this capability improves SCPS-TP’s estimate of round trip time, which can become distorted in error-prone environments.

- SCPS-TP Header Compression: this capability reduces protocol overhead by reducing the size of SCPS-TP headers.
- Selective Negative Acknowledgment: this capability improves SCPS-TP's error response by providing detailed information about missing or corrupted data.

We wished to determine the extent to which each of these capabilities affected performance at various bit-error rates. We also wished to determine if there were any significant interactions between the options that would restrict the ability of a user or program to pick the options individually.

We met the objectives stated above. The process of hosting the SCPS-TP protocol onto the STRV was a difficult one, primarily due to the limited availability of C-language development tools for the MIL-STD-1750A processor. The generally poor quality of development tools delayed our discovery and correction of two implementation errors. This rendered invalid the results of the first set of tests that we conducted. We were able to conduct a limited amount of retesting, which was used to confirm the results we gathered in the laboratory.

F4 SUMMARY OF RESULTS

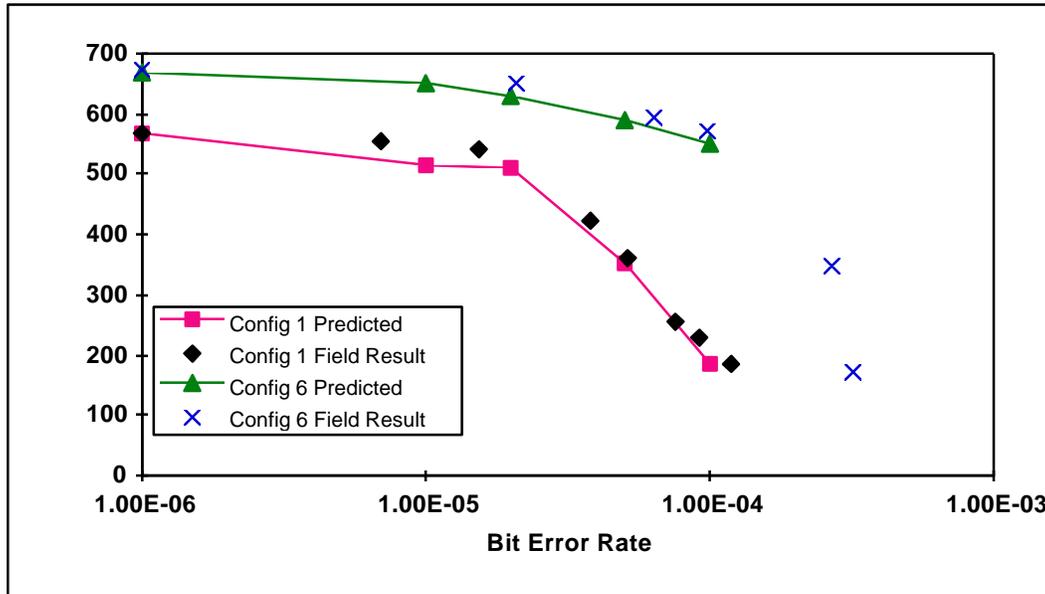
The experiment was conducted in three phases: we performed initial testing in the field, then we tested the protocols extensively in the laboratory, then we performed final testing in the field. The field results presented in this report reflect the results of the final field testing.

We made the following performance measurements in all tests: throughput, link utilization, and bit-efficiency. Throughput is a measure of the average rate at which the protocol can move user data, and is one of the most commonly used measurements of communication protocol performance. Link utilization is a measure of the ability of the protocol to "keep the pipe full." This ability is important in space communication, in which contact times may be limited. The protocol should not allow the link to be idle when data is waiting to be transmitted. Finally, bit-efficiency is a measure of the amount of protocol overhead required to transfer a user's data. The overhead includes protocol headers, acknowledgment traffic, and any retransmissions required to get the user data to its destination. Bit efficiency is important in spacecraft communications, because link capacity is generally a scarce resource.

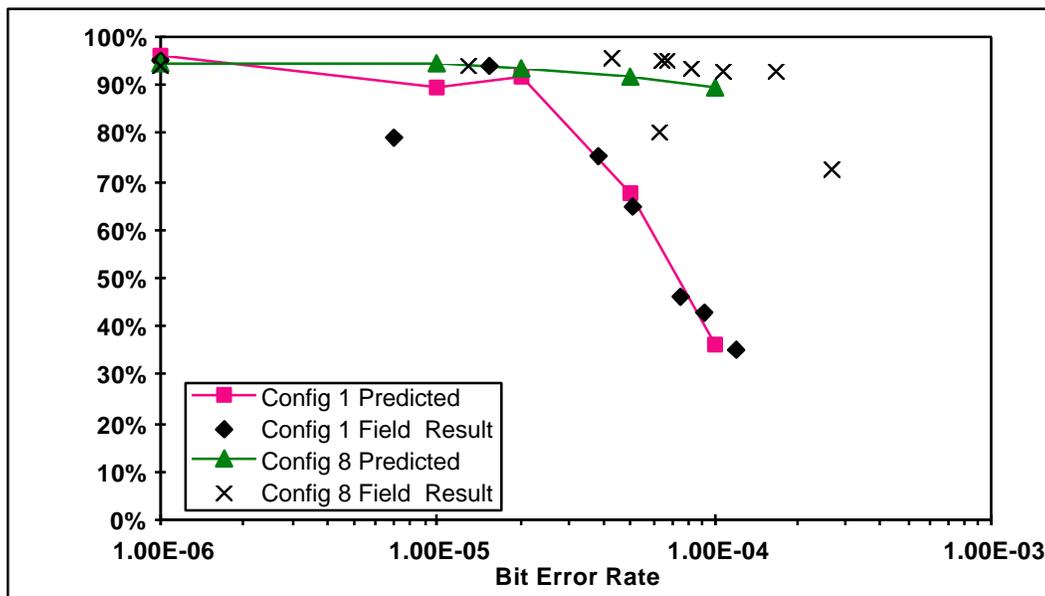
Eight protocol configurations were tested in the laboratory, and of those eight, five were tested in the field. The following graphs briefly summarize the field test results, with predictions based on the laboratory results shown for reference. The three graphs correspond to the three types of performance measures made: throughput, link utilization, and bit-efficiency. Each graph presents the results of the protocol configurations that performed the best and the worst in the field. The laboratory predictions are based on the mean response resulting from 10 tests at each of the following bit-error rates: 10^{-6} , 10^{-5} , 2×10^{-5} , 5×10^{-5} , and 10^{-4} .

The first graph presents the throughput results. Readers should bear in mind that the maximum possible throughput of a SCPS-TP connection is 768 bps, not including SCPS-TP protocol overhead. The graph shows that the best throughput was obtained by the configuration (Configuration 6) that enabled the Selective Negative Acknowledgment and SCPS-TP Header Compression capabilities, described above. The poorest throughput in the field resulted from the configuration (Configuration 1) that had none of the SCPS-TP capabilities enabled. (Note that

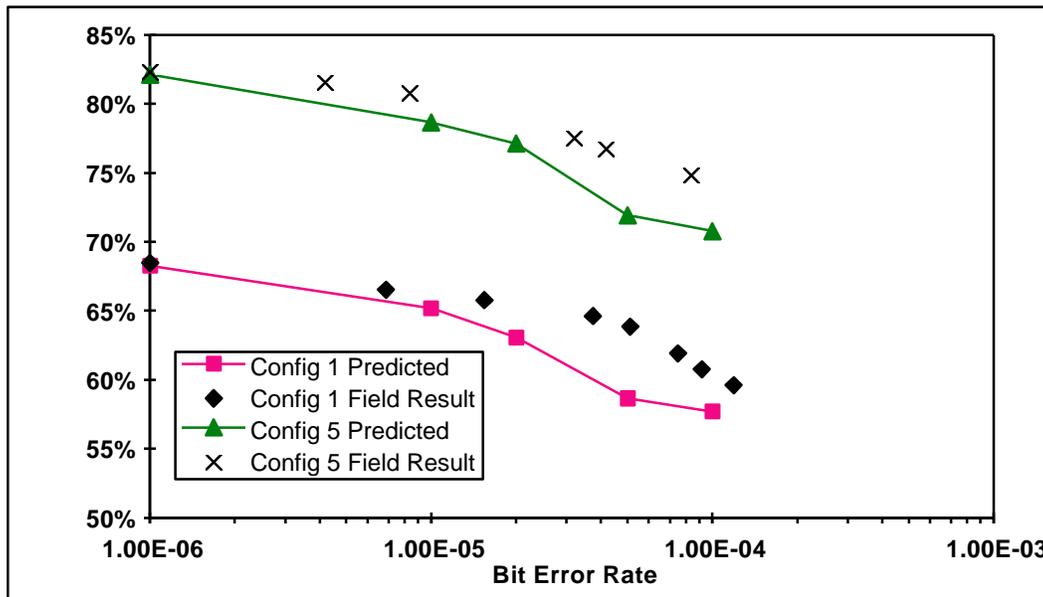
the laboratory results indicate that the configuration that enabled TCP Timestamps and none of the other capabilities would have shown lower throughput than Configuration 1, but this configuration was not tested in the field.)



The next graph presents link utilization results. The graph shows that the best link utilization was obtained by the configuration (Configuration 8) that enabled all of the capabilities under test. The poorest link utilization in the field resulted from the configuration (Configuration 1) that had none of the SCPS-TP capabilities enabled.



The final graph presents bit-efficiency results. The graph shows that the best bit-efficiency was obtained by the configuration (Configuration 5) that enabled SCPS-TP Header Compression. The poorest bit-efficiency in the field resulted from the configuration (Configuration 1) that had none of the SCPS-TP capabilities enabled. However, as with the throughput tests, the configuration with TCP Timestamps (only) enabled had worse bit-efficiency than Configuration 1 in the laboratory tests, but was not tested in the field. Note that the bit-efficiency results from the field tests tend to be higher than the results from the laboratory tests of corresponding configurations. This is due to an inherent difference between the field test environment and the laboratory test environment.



F5 Conclusions

The SCPS-TP protocol appears to be well-suited to the long-delay, potentially high bit-error rate environment of the STRV. All configurations were able to sustain connections at bit-error rates of 10^{-4} and yield throughput in excess of 130 bps (17% of maximum possible). The Selective Negative Acknowledgment (SNACK) capability was principally responsible for the ability of the SCPS-TP to operate well in high bit-error rate environments (the configuration with the SNACK capability enabled showed only a 15% drop from maximum throughput at a bit-error rate of approximately 10^{-4}). The SCPS-TP Header Compression accounted for an 18% increase in throughput over the configurations that did not use Header Compression at zero bit error rate.

The following conclusions derive from the laboratory testing and are confirmed by the flight test results:

- a) The SNACK capability significantly improves throughput at high bit-error rates, and has no negative effects on throughput at low bit-error rates.

- b) The TCP Timestamps capability has a negative effect on throughput at low bit-error rates. It has a strongly negative effect on bit-efficiency. When used in combination with SNACK, throughput is lower than when using SNACK alone. (The magnitude of the negative effect of TCP Timestamps on throughput is exaggerated by the small packet size imposed by the STRV. With larger packet sizes, this effect is mitigated.)
- c) The SCPS-TP Header Compression capability has a significant, positive effect on throughput at bit-error rates of 5×10^{-5} and below. Header Compression improves bit-efficiency at all bit-error rates. (The positive effect of Header Compression on throughput is exaggerated by the small packet size imposed by the STRV in the same manner that the negative effect of the TCP Timestamps is, above. As with TCP Timestamps, the effect of Header Compression on throughput will diminish as the packet size increases.)

The following conclusions are supported by the laboratory testing, but were neither confirmed nor refuted by the flight test results:

- a) The SNACK capability significantly improves link utilization at high bit-error rates, has no negative effects on link utilization at low bit-error rates, and has no impact on bit-efficiency.
- b) The TCP Timestamps capability has a moderately positive effect on link utilization. When used in combination with SNACK, link utilization is improved slightly.
- c) The SCPS-TP Header Compression capability has no effect on link utilization.

F6 Recommendations

The following recommendations are directed toward potential users of SCPS-TP and toward the sponsors of this effort.

- a) Push ahead in the effort to standardize SCPS-TP and deploy it in environments that have similar delay and error characteristics to the STRV environment.
- b) When using SCPS-TP in STRV-like environments, enable SNACK.

SNACK has no negative effects when errors are not present, and is primarily responsible for the protocol's ability to sustain relatively high throughputs at high bit-error rates.

- c) When using SCPS-TP in STRV-like environments, enable Header Compression.

The Header Compression capability reduced the size of SCPS-TP headers, improving throughput and bit-efficiency. These effects were particularly dramatic because the maximum packet size of the STRV was small. As the packet size increases, the positive effect of Header Compression will diminish.

- d) When using SCPS-TP in STRV-like environments, disable TCP Timestamps.

The TCP Timestamps capability reduced throughput at low bit-error rates, and provided no significant improvement in throughput at high bit-error rates when SNACK was in use. As with Header Compression, the negative effects of TCP Timestamps are exaggerated by the small packet sizes on STRV.

- e) Evolve the program of testing toward integrated tests.

Although there are still specific SCPS-TP capabilities to be tested, the focus of future tests should be integrated-stack testing. Tests of individual protocol capabilities can be conducted either as part of integrated-stack testing or as a small, focused portion of a larger test. The SCPS-NP, which has not as yet undergone flight testing, will probably benefit from more substantial, focused testing. However, this can still be conducted in the context of an overall test.