NATO STANDARDIZATION AGREEMENT
(STANAG)

TECHNICAL STANDARDS FOR HF RADIO LINK LAYER
AND APPLICATION SUPPORT PROTOCOLS

RELATED DOCUMENTS:

> See Section 6 (References).

AIM

The aim of this agreement is to define the functions and interfaces required for networked, error- free communication over HF radio channels, nominally for beyond-line-of-sight communications. This STANAG makes use of waveforms standardized in other STANAGs and standardized applications. This STANAG defines the necessary intermediate layers to provide a complete solution.

AGREEMENT

The participating nations agree to implement the protocols and interfaces defined in this STANAG (including mandatory Annexes) to provide communications over HF radio channels.

IMPLEMENTATION OF THE AGREEMENT

This STANAG is implemented by a nation when data communication on long-haul or short-haul HF radio channels complies with the characteristics detailed in this agreement.

DEFINITIONS

node | An implementation of the profile described in the main body of and mandatory annexes to this STANAG. The node is generally assumed to include the HF (modem and radio) and cryptographic equipment required for communications.

subnetwork | A collection of nodes. As a whole, a subnetwork provides a reliable networked data-transport service for external users or clients.

1. INTRODUCTION

This STANAG specifies protocols for data communication over HF radio, which will usually be used for beyond-line-of-sight communication. This STANAG describes a set of functions, segregated logically into layers, together with the interfaces, data formats, and procedures required for interoperability.   External standards and specifications are referenced and used where appropriate. The technical characteristics that are required to ensure interoperability and reliable system operation are described in the main body of and mandatory annexes to the document. Information-only annexes provide information on possible implementation of interfaces and subnetwork clients. The annexes also contain implementation advice based on experience during the development and deployment of the protocols.

This document is organised so that the main body gives an overview of the structure of the STANAG and the capabilities that should be realised when it is implemented. The details of the interfaces, data formats, and procedures are described in a number of mandatory and informational annexes.

2. SERVICES PROVIDED

This STANAG enables interoperability at the two major classes of interface: first, the "common air interface", describing how information is exchanged between nodes by radio; and second, the non-HF interfaces which allow external users or clients to interact with the subnetwork and with each other over the subnetwork. While physical interfaces are left up to the system implementer (e.g., Ethernet, FDDI, internal bus, or shared memory), the data formats (primitives) and procedures that make up the interface are specified in detail so that client applications can make use of the subnet.

2.1. Common Air Interface: Reliable Data Communications over HF Radio

Reliable data communications over HF radio is provided by using an ARQ data link protocol supported by the modern waveforms referenced in this STANAG. The STANAG is suitable for use with other waveforms with similar characteristics.

The data transfer sublayer defined in the profile supports automatic changes of the user data rate (that is, code rate) of the HF modem in response to changing channel conditions (adaptive data rate). This capability requires remote control of the HF modem and use of auto-baud waveforms such as STANAG 4539 or STANAG 5069. The profile is defined so nodes in which remote control of the modem, and hence adaptive data rate, is not available will interoperate with nodes which do have the capability.

The profile defined here supports frequency selection by use of Automatic Link Establishment (ALE), but does does not require that ALE is available.
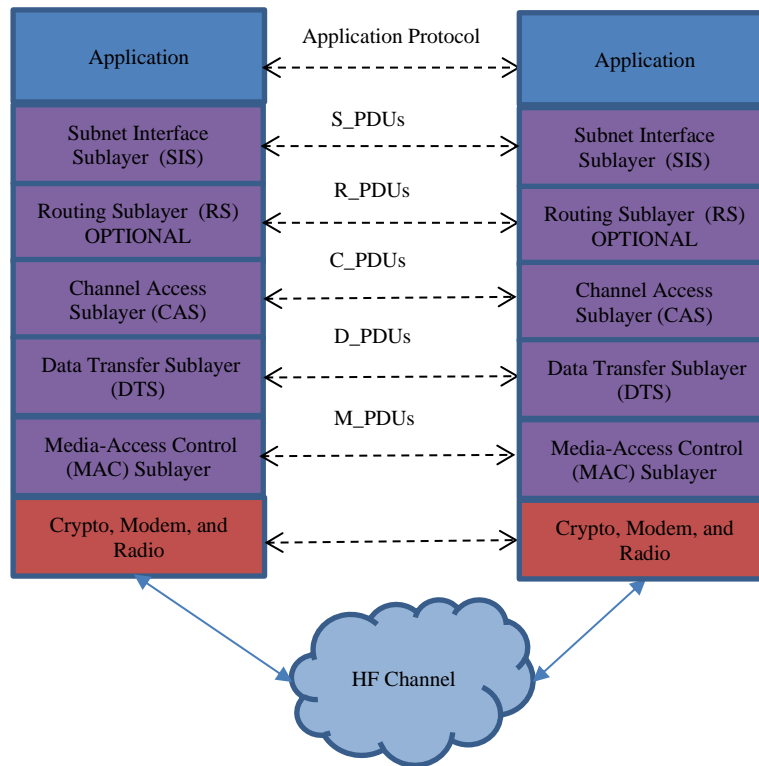
2.2.   <u>The STANAG 5066 Protocol Stack</u>

Figure 1: STANAG 5066 Layer Architecture

STANAG 5066 provides and interface between applications and the HF channel communications, comprising Crypto, Modem and Radio. The interface to the HF Channel is described in Section 3. Applications are described in Section 4.

STANAG 5066 has a layered architecture, with peer protocol defined with a set of PDUs and service interfaces specified between each layer.   The following sublayers are specified:

1. The **Subnetwork Interface Sublayer (SIS)** provides a service interface to applications using STANAG 5066. Annex A defines the service interface provided by SIS to applications. The primary service function is to transfer blocks of data (UNIDATA).   The SIS defines a peer protocol using S_PDUs.

2. The **Routing Sublayer (RS)** is an optional sublayer that is used in configurations where data needs to traverse an HF channel multiple times.   This is needed to support Wireless Token Ring Protocol (WTRP) configurations with partial connectivity. The RS is specified in Annex R.

3. The **Channel Access Sublayer (CAS)** controls communication between peers by managing "physical links" which control ARQ communication between peers.   The CAS can provide simultaneous access to multiple peers or constrain operation to one active peer at a time.   The CAS is specified in Annex B.

4. The **Data Transfer Sublayer (DTS)** provides data transfer with one or more peers. The DTS provides a reliable (ARQ) data link service, as well as best-effort (non-ARQ) service for broadcast, multicast and communication to peers in EMCON. The DTS supports data rate selection for STANAG 4539

and STANAG 5069 waveforms. The DTS is specified in Annex C.

5. The **Media-Access-Control (MAC) Sublayer** provides mechanisms for enhanced media-access control capability for HF data communication in multi-node networks. Whereas the Channel-Access Sublayer provides pairwise logical link control mechanisms to establish a point-to-point link (or set of multiple, independent point-to-point links) for data communication, the Media-Access- Control Sublayer introduces modes for enhanced media-access control capability for HF data communication in multi-node networks, and the prescribed method in which they may be used with other STANAG 5066 capabilities. These optional channel-access modes extend or modify, but do not replace, the pairwise logical link control mechanisms defined for the Channel Access Sublayer. General requirements for the Media-Access-Control Sublayer are defined in Annex J of this STANAG, with requirements on each of the defined protocols for media-access- control provided in Annexe K for Carrier-Sense-Multiple Access and Annex L for Wireless-Token- Ring Protocol (WTRP).   Annexes may be added in future versions of this STANAG for other techniques, such as Adaptive Time Division Multiple Access.

3. STANAG 5066 Interfaces to the HF Channel.

STANAG 5066 defines three ways to interface to the HF channel, which are described in this section. These interfaces can all be used with the waveforms defined in STANAG 4285, STANAG 4529, STANAG 4539 and STANAG 5066. It is anticipated that STANAG 5066 can be used with other waveforms with similar characteristics to these.

3.1. <u>Direct Interface to Modem</u>

**Figure 2: STANAG 5066 Direct Interface to Modem**

The STANAG 5066 stack **may** interface directly to a modem, with a stack as shown in Figure 2. Data interface between STANAG 5066 stack and Modem **may** use the synchronous serial interface specified in Annex D. This data interface **may** use other communication mechanisms such as the TCP data interface specified in Annex A of MIL-STD-188-110D.

3.2. Interface to Synchronous Serial Crypto Equipment

Figure 3: STANAG 5066 Interface to Synchronous Serial Crypto Equipment

The STANAG 5066 stack can be interfaced to Synchronous Serial Crypto Equipment, such as BID-950, KG-84C and KIV-7, as shown in Figure 3. Data interface between STANAG 5066 and Crypto Equipment **shall** conform to Annex D of this specification.

3.3. Interface using AES and other Encryption Algorithms

Figure 4: STANAG 5066 Interface using AES and other Encryption Algorithms

The STANAG 5066 stack can be extended downwards to include encryption in a TRANSEC sublayer following the specification of Annex T of this specification. Annex T specifies use of the Advanced Encryption Standard (AES) and provides a general framework for use of alternative encryption algorithms.

Data interface between the TRANSEC layer and Modem **may** use the synchronous serial interface specified in Annex D. This data interface **may** use other communication mechanisms such as the TCP data interface specified in Annex A of MIL-STD-188-110D.

3.4. Modem Control and Monitoring

Figure 5: STANAG 5066 Modem and ALE Control

STANAG 5066 allows for, but does not require control communication with the modem. Monitoring and control of the modem is performed by the DTS, as shown in Figure 5. The following requirements are identified:
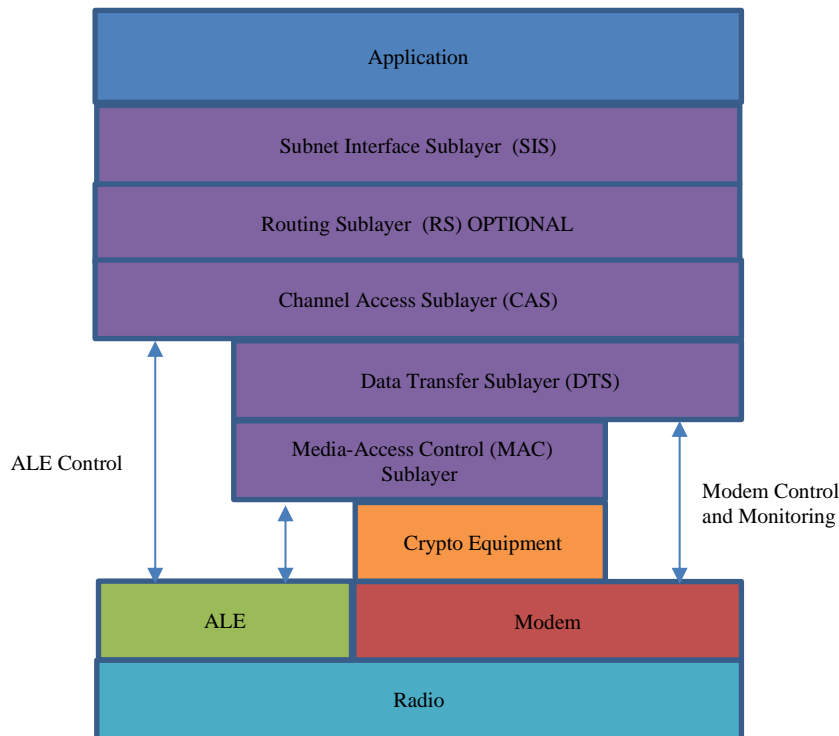
1. To achieve Data Rate Selection, as described in Annex C, the sender needs to control the modems sending speed and associated parameters such as interleaver.

2. To provide optimum recommendations on transmission parameters, the receiver needs to be able to monitor SNR and other receive characteristics.

3. To determine switching time when no EOT is received, performance of Annex K and Annex L can be optimized by monitoring end of transmission at modem level.

In addition to these requirements and implementation may use additional modem control, configuration and monitoring capabilities to improve system performance and management.

3.1. Use of Automatic Link Establishment (ALE)

Automatic Link Establishment (ALE) enables selection of best frequency. For STANAG 5069, 4G ALE also allows selection of transmit and receive bandwidth. ALE mechanisms are specified in MIL-STD-188-141D, providing 2G (synchronous), 3G (asynchronous) and 4G mechanisms.   STANAG 4538 defines 3G mechanisms.

ALE sits at the same level as the modem in the protocol architecture. ALE and Modem share use of the radio, with only one of them having access at any moment.   ALE and Modem will co-ordinate sharing of the modem.

STANAG 5066 specifies two uses of ALE, as shown in Figure 5:

1. ALE used by MAC Sublayer, as specified in Annex J. In this mode, the MAC layer negotiates a frequency for use by all nodes on the subnetwork.

2. ALE used by the CAS, as specified in Annex B. In this mode, ALE is used to connect to one peer node or to a subset of nodes on the network. This will lead to the SIS queueing traffic for other nodes, as specified in Annex A.

ALE **may** be used in a third implicit mode on a two-node network. In this configuration, ALE can be used to set up a link without interaction with STANAG 5066, as there is only one destination.

3.2. Crypto Bypass

It can be seen in Figure 5 that all user data is passed through the Crypto. In order to achieve use of ALE and variable data rate, STANAG 5066 needs to communicate with components on the other side of the Crypto. This communication is called Crypto Bypass.

When Crypto Bypass is used, additional controls and checks **may** be needed on the information flow. Mechanisms to achieve this are outside of the scope of this STANAG.

4. APPLICATION SUPPORT

The SIS service provides a simple service interface to support applications, as defined in Annex A.

4.1. Applications Supported

Each application supported by STANAG 5066 will use a SAP (Subnet Access Point). STANAG 5066 can support up to 16 applications sharing a channel, each with a different SAP ID (0-15).

Annex F specifies SAP ID assignment for a number of applications. Where one of these applications is used, the assigned SAP ID **shall** be used. If an application listed in Annex F is not being used, a deployment **may** use its assigned SAP ID for a different application.

STANAG 5066 fully specifies one application, specifies supporting protocol for other applications, and references external applications. This includes:

1. HF Operator Chat. A simple operator chat protocol specified in Annex O.

2. Support for ACP 127 legacy formal military messaging using the Character Oriented Stream Service (COSS), specified in Annex P.

3. Support for Multicast Transfer using ACP 142, specified in Annex Q. This enables two services over ACP 142:

    a. STANAG 4406 Annex E, to provide formal military messaging.

    b. Multicast Email (MULE) specified in RFC 8494, to provide email or formal military messaging using RFC 6477.

4. Compressed File Transfer Protocol (CFTP) which can be used to provide basic email services, specified in Annex V.

5. Chat and Presence using XMPP, specified in XEP-0365.

6. IP Client, that enables an IP subnet to be operated over STANAG 5066. This provides support for some IP services, such as ICMP Ping. IP Client is specified in Annex U.

4.2. <u>Enabling Applications to Share a STANAG 5066 Stack</u>

Figure 6: SIS Access Protocol Model

It is desirable to have multiple applications share a single STANAG 5066 Stack, so that a common stack implementation can be used and that data from the different applications can be multiplexed over STANAG 5066 communication. This is achieved by use of the SIS Access Protocol, illustrated in Figure 6 and specified in Annex S. The SIS Access Protocol operates locally to a node and has no end to end implications.

5. List of Annexes

Annex A:	Subnetwork Interface Sublayer (mandatory)
Annex B:	Channel Access Sublayer (mandatory)
Annex C:	Data Transfer Sublayer (mandatory)
Annex D:	Interface between Data Transfer Sublayer and Communications Equipment (mandatory)
Annex E:	Absent
Annex F:	SAP Assignment (mandatory)
Annex G:	Absent
Annex H:	Absent
Annex I:	Absent
Annex J:	General Requirements for Enhanced Media-Access-Control (MAC) Capabilities in Multi-Node STANAG 5066 Networks
Annex K:	High-Frequency Carrier-Sense Multiple-Access (CSMA) Protocols
Annex L:	High-Frequency Wireless Token-Ring-Protocol (WTRP) Requirements
Annex M:	Reserved
Annex N:	Guidance on Address Management in STANAG 5066 Networks
Annex O:	HF Operator Chat
Annex P:	ACP 127 & Character-Oriented Serial Stream
Annex Q:	ACP 142
Annex R:	Routing Sublayer
Annex S:	SIS Access Protocol (mandatory)
Annex T:	STANAG 5066 TRANSEC Crypto Sublayer using AES and other Protocols
Annex U:	IP Client
Annex V:	Compressed File Transfer Protocol

Annexes not marked as "mandatory" are "information only".

Annexes marked "Absent" were present in previous editions of STANAG 5066. Annex M is marked "reserved" as a provisional allocation has been made.

6. References

1.   Clark, D., and N. Karavassillis, "Open Systems for Radio Communications: A Subnet Architecture for Data Transmission over HF Radio", TM-937, May 1998

2.   NATO STANDARDIZATION AGREEMENT 4203 — Technical Standards for Single Channel HF Radio Equipment

3.   NATO STANDARDIZATION AGREEMENT 4285— Characteristics of 1200/ 2400/ 3600 bps single tone modulators/demodulators for HF radio links

4.   NATO STANDARDIZATION AGREEMENT 4406 Annex E — Military Message Handling System, Annexe E: Tactical MMHS Protocol and Profile Solution, March 2005

5.   NATO STANDARDIZATION AGREEMENT 4529 — Characteristics of Single-Tone Modulators/Demodulators for Maritime HF Radio Links with 1240 Hz bandwidth

6.   NATO STANDARDIZATION AGREEMENT 4538 — Technical Standards for an Automatic Radio Control System (ARCS) for HF Communications Links

7.   NATO STANDARDIZATION AGREEMENT 4539 — Technical Standards for Non-Hopping HF Communications Waveforms

8.   NATO STANDARDIZATION AGREEMENT 5069 — Technical??

9.   DEPARTMENT OF DEFENSE INTERFACE STANDARD 188-110D — Interoperability and Performance Standards for Data Modems, MIL-STD-188-110D, 29 December 2017

10.   DEPARTMENT OF DEFENSE INTERFACE STANDARD 188-141D — Interoperability and Performance Standards for Medium and High Frequency Radio Systems, MIL-STD-188-141D, 29 December 2017

11.   Combined Communications-Electronics Board ACP127(G), "Tape Relay Procedures", November 1988

12.   Combined Communications-Electronics Board ACP142A, "P_MUL - A Protocol for Reliable Multicast Messaging in Bandwidth Constrained and Delayed Acknowledgement (EMCON) Environments", October 2008

13.   CCITT V41 - Code-independent error-control system, November 1988

14.   CCITT V.42 - Error-correcting procedures for DCEs using asynchronous-to-synchronous conversion, March 2002

15.   National Institute of Standards and Technology (NIST), "Advanced Encryption Standard (AES)", 1988

16.   REQUEST FOR COMMENTS 791 — Postel, J., "Internet Protocol", September 1981.

17.   REQUEST FOR COMMENTS 6477 — Melnikov, A; Lunt, G., "Registration of Military Message Handling System (MMHS) Header Fields for Use in Internet Mail", January 2012

18.   REQUEST FOR COMMENTS 8200 — Dering, S.; Hinden, R., "Internet Protocol, Version 6 (IPv6) Specification". July 2017

19.   REQUEST FOR COMMENTS 8494 — Wilson, D.; Melnikov, A., "Multicast Email (MULE) over Allied Communications Publication (ACP) 142", November 2018

20.   XMPP Standards Foundation XEP-0365, "Server to Server communication over STANAG 5066 ARQ", Kille, S., July 2018

21.   Miller, T., and P. Reynolds, "Experience with Approved Cryptographic Equipment in HF ARQ Systems", NC3A TN 638, NATO CONFIDENTIAL, November 1996

22.   Mustafa Ergen, Duke Lee et. al., "Wireless Token Ring Protocol", University of California, Berkeley, CA 94720, USA.

7. Changes in Edition 4

Most changes since Edition 3 are in the annexes. Each annex has a section at the end listing changes in Edition 4. The following changes are made in the main document:

1.   New Title that better reflects purpose.

2.   Updated overall architecture to reflect new annexes.

3.   General update of the text to improve clarity.

4.   Added section to summarize use of ALE.

The following table lists the annexes and summarizes the major changes since Edition 3.

Annex	Edition 4 Title	Key Changes
A	Subnetwork Interface Sublayer	Removal of Hard Links, Expedited, and Rank, which are not needed. Moved access protocol to Annex S.
B	Channel Access Sublayer	Added Edition 4 negotiation and support for ALE.
C	Data Transfer Sublayer	Extended frame sequence number to support higher speeds. Provided D_PDU extensibility.
D	Interface between Data Transfer Sublayer and Communications Equipment	No significant change.
E	Absent	Modem control interface that was out of date and not useful.
F	SAP Assignment	Most content moved to annexes O-Q, S, U and V. Renamed to reflect core content. Some new assignments made
G	Absent	Requirements on 2400 bps modems, now handled in Annex D.
H	Absent	Implementation notes. Information still valid folded into Annex D and other places as appropriate.
I	Absent	Frequency change procedure that is no longer needed.
J	General Requirements for Enhanced Media-Access-Control (MAC) Capabilities in Multi-Node STANAG 5066 Networks	No significant change.
K	High-Frequency Carrier-Sense Multiple-Access (CSMA) Protocols	Added "slotted" support to optimize performance on small networks.
L	High-Frequency Wireless Token-Ring-Protocol (WTRP) Requirements	Model unchanged. Significant update of protocol.
M	Reserved	
N	Guidance on Address Management in STANAG 5066 Networks	No significant change.
O	HF Operator Chat	Extracted from Annex F. No significant change.
P	ACP 127 & Character-Oriented Serial Stream	Extracted from Annex F. No significant change.
Q	ACP 142	Extracted from Annex F.

			Description changed (no protocol change) to support MULE as well as STANAG 4406 Annex E.
R	Routing Sublayer		New optional sublayer
S	SIS Access Protocol		Extracted from Annexes A and F. No significant change.
T	STANAG 5066 TRANSEC Crypto Sublayer using AES and other Protocols		New optional layer to provide alternate TRANSEC to Annex D.
U	IP Client		Extracted from Annex F. Added IPv6 support.
V	Compressed File Transfer Protocol		Extracted from Annex F. No significant change.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

Annex A: Subnetwork Interface Sublayer
(mandatory)
DRAFT 1.0 of 9th April 2020

This annex defines the interface between the users of the HF subnetwork and the computer information system through which the user accesses the subnetwork.

A.1   Subnetwork Service Definition

The Subnetwork Service can be used to support multiple applications and multiplexes multiple applications over a single channel. The service provided by the server is application independent and common to all clients irrespective of the task they may perform.

Clients are associated to the Subnetwork Interface Sublayer by Subnetwork Access Points (SAPs). There can be multiple clients simultaneously using the Subnetwork Interface Sublayer. Each SAP is identified by its SAP Identifier (SAP ID). The SAP ID is a number in the range 0-15; hence there can be a maximum of 16 clients using the Subnetwork Interface Sublayer of a single node. SAPs are equivalent to the "ports" of the TCP and UDP protocols.

Clients **may** connect to the Subnet Interface Service using the SIS Access Protocol specified in STANAG 5066 Annex S. Clients **may** use SIS layer directly or connect with a different protocol. Clients are responsible for segmenting larger messages into User Protocol Data Units (U_PDUs).

The Subnetwork Interface Sublayer treats all clients connected to it in the same manner irrespective of the application performed by these clients.

A.1.1   Initiating Data Exchange Sessions

The Subnetwork Interface Sublayer is responsible for initiating the establishment and termination of Sessions with its peers at remote nodes. There are two types of sessions:

   A.2.1.1 Soft Link Data Exchange Session, which require the making of a point-to-point physical link with a specified remote node.

   A.2.1.2 Broadcast Data Exchange Session

Clients for the HF Subnetwork services **may** interleave requests for the various session types in accordance with the capabilities of this standard. Support for only one session type, e.g., restriction to support only a Broadcast Data Exchange Session, **may** be established as part of the local (implementation-dependent) subnetwork management function.

A.1.2   Soft Link Data Exchange Session

The establishment of a Soft Link Data Exchange Session **shall** [(1)] be initiated unilaterally by the Subnetwork Interface Sublayer which has queued data requiring reliable delivery (i.e., queued ARQ U_PDUs). Soft Links **may** also be initiated for Non-ARQ traffic for use over ALE links, when requested by the Channel Access Sublayer.

The Subnetwork Interface Sublayer **shall** [(2)] initiate Soft Link Data Exchange Sessions as needed, following the

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

procedure described in Section A.3.2.1.

When all data has been transmitted to a node with which a Soft Link Data Exchange Session has been established, the Subnetwork Interface Sublayer **shall** [3] terminate the Soft Link Data Exchange Session after a configurable and implementation-dependent time-out period in accordance with the protocol specified in Section A.3.2.1.3.

Termination of the Soft Link Data Exchange Session **shall** [4] be in accordance with the procedure specified in Section A.3.2.1.3. The time out period may be zero. The time out period allows for the possibility of newly arriving U_PDUs being serviced by an existing Soft Link Data Exchange Session prior to its termination.

In order to provide "balanced" servicing of the queued U_PDUs, a Soft Link Data Exchange Session **shall** [5] not be maintained for a period which exceeds a specified maximum time if U_PDUs of appropriate priorities are queued for different node(s).

The specified maximum time out period **shall** [6] be a configurable parameter for the protocol implementation. The specific values of the parameters governing the establishment and termination of Soft Link Data Exchange Sessions (e.g. time-out periods etc.) are chosen in the context of a particular configuration (i.e. size of network, etc).

A.1.3   Broadcast Data Exchange Session

The third type of data exchange session is the Broadcast Data Exchange Session. The subnetwork **shall** [1] service only clients with service requirements for non-ARQ U_PDUs during a Broadcast Data Exchange Session. [Note: Clients with service requirements for non-ARQ U_PDUs may be serviced during other session types, however, in accordance with the session's service characteristics.] A Broadcast Data Exchange Session can be initiated and terminated by a management process, e.g., a local or network administrator management client.

The procedures that initiate and terminate broadcast data exchange sessions **shall** [2] be as specified in Annex C.

A node configured to be a broadcast-only node **shall** [3] use a "permanent" Broadcast Data Exchange Session during which the Subnetwork Interface Sublayer **shall** [4] service no ARQ Data U_PDUs. Alternatively the Subnetwork Interface Sublayer can unilaterally initiate and terminate Broadcast Data Exchange Sessions.

A.2   Subnet Service Specification

The Service Interface to the Subnetwork Interface Sublayer provides the interface primitives listed in Table A-1 and defined in the following subsections. The names of these primitives are prefixed with an "S_" to indicate that they are exchanged across the interface between the subnetwork interface sublayer and the application using the service interface. This table is intended to provide a general guide and overview to the primitives. For detailed specification of the primitives, the later sections of this Annex **shall** [1] apply.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

Table A-1. Primitives Exchanged with Clients

CLIENT -> SUBNETWORK INTERFACE	SUBNETWORK INTERFACE -> CLIENT
S_BIND_REQUEST (Service Type, SAP ID)	S_BIND_ACCEPTED (SAP ID, MTU)
	S_BIND_REJECTED (Reason)
S_UNBIND_REQUEST ( )	S_UNBIND_INDICATION (Reason)
S_UNIDATA_REQUEST (Destination Node Address, Destination SAP ID, Priority, TimeToLive, Delivery Mode, U_PDU)	S_UNIDATA_REQUEST_CONFIRM (Destination Node Address, Destination SAP ID, Size of confirmed U_PDU, U_PDU)
	S_UNIDATA_REQUEST_REJECTED (Reason, Destination Node Address, Destination SAP ID, Size of Rejected U_PDU, U_PDU)
	S_UNIDATA_INDICATION (Source Node Address, Source SAP ID, Destination Node Address, Destination SAP ID, Priority, Transmission Mode, *transmission-mode conditional parameters* , U_PDU)

A.2.1   Management and Flow Control

It is anticipated that the service interface between the SIS layer and application will contain:

1. Flow Control to control flow of data between SIS and Application.
2. Management information, to provide the application with additional information

These functions are not standardized in this Annex, and **may** be chosen to support the application. These functions are local and do not impact end to end interoperability.

Annex S specifies a management and flow control primitives that are appropriate to use with this annex.

A.2.2   Content Specification and Use of Primitives

The content specification and use of the Subnetwork Interface Sublayer primitives **shall** [1] be as specified in the following subsections.

A.2.2.1   S_BIND_REQUEST Primitive

Name :
    S_BIND_REQUEST ( )
Arguments :
    3. SAP ID,
    4. Default Service Type
Direction :
    Client -> Subnetwork Interface
Description :
    The S_BIND_REQUEST primitive **shall** [1] be issued by a new client when it first connects to the subnetwork. Unless this primitive is issued the client can not be serviced. With this primitive the client uniquely identifies and declares that it is "on-line" and ready to be serviced by the subnetwork.

    The first argument of this primitive **shall** [2] be the "*SAP ID*" which the client wishes to be assigned. The SAP ID **shall** [3] be node-level unique, i.e. not assigned to another client connected to the Subnetwork Interface Sublayer for a given node.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The last argument of this primitive **shall** [(7)] be "*Service Type*" and identifies the default type of service requested by the client and specified in Section A.2.2.6. The *Service Type* argument **shall** [(8)] apply to all data units submitted by the client unless explicitly overridden by client request when submitting a U_PDU to the subnetwork.

A.2.2.2   S_UNBIND_REQUEST Primitive

Name :
          S_UNBIND_REQUEST ( )
Arguments :
          NONE
Direction :
          Client -> Subnetwork Interface ( )

Description :
          The S_UNBIND_REQUEST primitive **shall** [(1)] be issued by a client in order to declare itself "off-line". The Subnetwork Interface Sublayer **shall** [(2)] release the SAP ID allocated to the client from which it receives the S_UNBIND_REQUEST and the SAP_ID allocated to this client **shall** [(3)] then be available for allocation to another client that may request it.

          A client that went off-line by issuing the S_UNBIND_REQUEST primitive can come on-line again by issuing a new S_BIND_REQUEST.

          A client can also go off-line by physically disconnecting itself (e.g. powering down the computer which runs the client program) or disconnecting the physical cable (RS232, Ethernet, etc.) which may connect the client to the node.

          The Subnetwork Interface Sublayer can sense whether a client is physically disconnected in order to unilaterally declare this client as off-line; the S_KEEP_ALIVE primitive specified in Section **Error! Reference source not found.** provides this capability, though other implementation-dependent methods may be used in addition to this primitive.
                    [Note: The omission of SAP ID as an argument in this and other primitives implies a requirement on the stack supporting this connection to associate a SAP ID with a lower level connection (i.e., socket) and maintain this association.]

A.2.2.3   S_BIND_ACCEPTED Primitive

Name :
          S_BIND_ACCEPTED ( )
Arguments :
          1. SAP ID
          2. Maximum Transmission Unit (MTU)
Direction :
          Subnetwork Interface -> Client
Description :
          The S_BIND_ACCEPTED primitive **shall** [(1)] be issued by the Subnetwork Interface Sublayer as a positive response to a client's S_BIND_REQUEST.

          The *SAP ID* argument of the S_BIND_ACCEPTED primitive **shall** [(2)] be the SAP ID assigned to the client and **shall** [(3)] be equal to the *SAP ID* argument of the S_BIND_REQUEST to which this primitive

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

is a response.

The *MTU* argument **shall** [(4)] be used by the subnetwork interface sublayer to inform the client of the maximum size U_PDU (in bytes or octets) which will be accepted as an argument of the S_UNIDATA_REQUEST primitive. S_UNIDATA_REQUEST primitives containing U_PDUs larger than the MTU **shall** [(5)] be rejected by the subnetwork interface. Note that this restriction applies only to U_PDUs received through the subnetwork interface. U_PDUs which are received from the lower HF sublayers (i.e., received by radio) **shall** [(6)] be delivered to clients regardless of size.

For general-purpose nodes, the MTU value **shall** [(7)] be 2048 bytes. For broadcast-only nodes, the MTU **shall** [(8)] be configurable by the implementation up to a maximum that **shall** [(9)] not exceed 4096 bytes.

A.2.2.4   S_BIND_REJECTED Primitive

Name :
          S_BIND_REJECTED ( )
Arguments :
          1.   Reason
Direction :
          Subnetwork Interface -> Client
Description :
          The S_BIND_REJECTED primitive **shall** [(1)] be issued by the Subnetwork Interface Sublayer as a negative response to a client's S_BIND_REQUEST. If certain conditions are not met then the Subnetwork Interface Sublayer rejects the client's request.

          The *Reason* argument of the S_BIND_REJECTED primitive **shall** [(2)] specify the reason why the client's request was rejected. Valid *Reason* values **shall** [(3)] be as specified in the table below.

Reason	Value
Not Enough Resources	1
Invalid SAP ID	2
SAP ID already allocated	3
ARQ Mode unsupportable during Broadcast Session	4

     The value assigned to each reason **shall** be used to represent the reason in SIS Access Protocol (Annex S).

A.2.2.5   S_UNBIND_INDICATION Primitive

Name :
          S_UNBIND_INDICATION ( )
Arguments :
          1.   Reason
Direction :
          Subnetwork Interface->Client
Description :
          The S_UNBIND_INDICATION primitive **shall** [(1)] be issued by the Subnetwork Interface Sublayer to unilaterally declare a client as off-line. If the client wants to come on-line again, it must issue a new a S_BIND_REQUEST primitive as specified in Section A.2.1.1.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The S_UNBIND_INDICATION primitive provides a means for the Subnetwork Interface Sublayer to manage the clients connected to it.

The *Reason* argument of the S_UNBIND_INDICATION primitive **shall** [(2)] specify why the client was declared off-line.

Reason	Value
Reserved	1
Inactivity (failure to respond to "Keep alive")	2
Too many invalid primitives	3
Reserved	4
ARQ Mode Unsupportable during Broadcast Session	5
SIS Service Terminating	6

The value assigned to each reason **shall** be used to represent the reason in SIS Access Protocol (Annex S).

A.2.2.6   S_UNIDATA_REQUEST Primitive

Name :
        S_UNIDATA_REQUEST( )
Arguments :
        1.   Priority
        2.   Destination SAP ID
        3.   Destination Node Address
        4.   Delivery Mode (Service Type)
        5.   TimeToLive (TTL)
        6.   Size of U_PDU
        7.   U_PDU (User Protocol Data Unit)
Direction :
        Client->Subnet Interface
Description :

The S_UNIDATA_REQUEST primitive **shall** [(1)] be used by connected clients to submit a U_PDU to the HF subnetwork for delivery to a receiving client.

The argument *Priority* **shall** [(2)] represent the priority of the U_PDU. The U_PDU priority **shall** [(5)] take a value in the range 0-15, with 0 being the lowest priority and 15 the highest. The processing by HF protocol sublayers **shall** [(6)] make a "best effort" to give precedence to high priority U_PDUs over lower priority U_PDUs which are queued in the system.

The argument *Destination SAP ID* **shall** [(3)] specify the SAP ID of the receiving client. Note that as all nodes will have uniquely specified SAP IDs for clients, the Destination SAP ID distinguishes the destination client from the other clients bound to the destination node.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The argument *Destination Node Address* **shall** [4] specify the HF subnetwork address of the physical HF node to which the receiving client is bound.

The argument *Delivery Mode* (Service Type) **shall** [5] be comprised of the four elements specified here. This argument can be given the value of "DEFAULT" which means that the delivery mode associated with the U_PDU will be the Default Service Type specified by the client during "binding" (i.e., the value DEFAULT is equal to the *Service Type* argument of client's original S_BIND_REQUEST). Values other than DEFAULT for the *Delivery Mode* can be used to override these default values. Service Type comprises the following elements:

1. *Transmission Mode for the Service.* --- ARQ or Non-ARQ Transmission Mode **shall**[6] be specified, with one of the Non-ARQ submodes if Non-ARQ was requested. Non-ARQ transmission has submodes such as: *Error-Free-Only* delivery to destination client, *Delivery-with-Errors* which provides delivery to destination client even with *some* errors.

2. *Data Delivery Confirmation for the Service* --- The client **may** request none, one or both of the Data Delivery Confirmation modes for the service. There are two types of data delivery confirmation:
   - Node-to-Node Delivery Confirmation
   - Client-to-Client Delivery Confirmation

   The client can request explicit confirmation, i.e, Node-to-Node or Client-to-Client, from the Subnetwork to provide indication that its U_PDUs have been properly delivered to their destination. Explicit delivery confirmation **shall** [9] be requested only in combination with ARQ delivery.

   [Note: The Node-to-Node Delivery Confirmation does not require any explicit peer-to-peer communication between the Subnetwork Interface Sublayers and hence it does not introduce extra overhead. It simply uses the ACK (ARQ) confirmation provided by the Data Transfer Sublayer. Client-to-Client Delivery Confirmation requires explicit peer-to-peer communication between the Sublayers and therefore introduces overhead. It should be used only when it is absolutely critical for the client to know whether or not its data was delivered to the destination client (which may, for instance, be disconnected).]

   NOTE: This service definition allows for both types of delivery confirmation to be requested. Annex S (Edition 4) only allows for one type to be requested.

3. *Order of delivery of any U_PDU to the receiving client.* --- A client **shall** [10] request that its U_PDUs are delivered to the destination client "in-order" (as they are submitted) or in the order they are received by the destination node.

   "in order" **shall not** be requested for the Non-ARQ service.

   NOTE: Use of "in-order" is implemented in the DTS and this leads to interaction between different clients using "in-order" and can lead to unnecessary data loss. It is recommended to avoid use of "in-order" and to handle ordering within the client application.

4. *Minimum Number of Retransmissions* --- This argument **shall** [12] be valid if and only if the Transmission Mode is a Non-ARQ type. If the Transmission Mode is a Non-ARQ type, then the subnetwork **shall** [13] retransmit each U_PDU the number of times specified by this service. If this is not specified, the U_PDU is sent only once.
   [Note: In non-ARQ Mode, automatic retransmission a minimum number of times may be used to

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

improve the reliability of broadcast transmissions where a return link from the receiver is unavailable for explicit retransmission requests.]

The argument *TimeToLive (TTL)* **shall** [6] specify the maximum amount of time the submitted U_PDU is allowed to stay in the HF Subnetwork before it is delivered to its destination. If the TTL is exceeded the U_PDU **shall** [7] be discarded. A TTL value of 0 **shall** [8] define an infinite TTL**,** i.e. the subnetwork should try *forever* to deliver the U_PDU.

The subnetwork **shall** [9] have a default maximum TTL. The default maximum TTL **shall** [10] be configurable as an implementation-dependent value. As soon as the Subnetwork Interface Sublayer accepts a S_UNIDATA_REQUEST primitive, it **shall** [11] immediately calculate its *TimeToDie (TTD)* by adding the specified TTL (or the default maximum value if the specified TTL is equal to 0) to the current Time of Day. The TTD attribute of a U_PDU **shall** [12] accompany it during its transit within the subnetwork. [Note that the TTD is an absolute time while the TTL is a time interval relative to the instant of the U_PDU submission.]

The *Size of U_PDU* argument **shall** [13] be the size of the U_PDU that is included in this S_UNIDATA_REQUEST Primitive.

The final argument, *U_PDU,* **shall** [14] be the actual Data Unit submitted by the client to the HF Subnetwork.

A.2.2.7   S_UNIDATA_REQUEST_CONFIRM Primitive

Name :
S_UNIDATA_REQUEST_CONFIRM
Arguments :
1. Destination Node Address
2. Destination SAP ID
3. Size of Confirmed U_PDU
4. U_PDU (User Protocol Data Unit or part of it)
Direction :
Subnetwork Interface->Client

Description :
The S_UNIDATA_REQUEST_CONFIRM primitive **shall** [1] be issued by the Subnetwork Interface Sublayer to acknowledge the successful delivery of a S_UNIDATA_REQUEST submitted by the client.

This primitive **shall** [2] be issued only if the client has requested Data Delivery Confirmation (either during binding or for this particular data unit).

The *Destination Node Address* argument in the S_UNIDATA_REQUEST_CONFIRM Primitive **shall** have the same meaning and be equal in value to the *Destination Node Address* argument of the S_UNIDATA_REQUEST Primitive for which the S_UNIDATA_REQUEST_ CONFIRM Primitive is the response.

The *Destination SAP_ID* argument in the S_UNIDATA_REQUEST_ CONFIRM Primitive **shall** [4] have the same meaning and be equal in value to the *Destination SAP_ID* argument of the S_UNIDATA_REQUEST Primitive for which the S_UNIDATA_REQUEST_CONFIRM Primitive is the response.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The *Size of Confirmed U_PDU* argument **shall** [5] be the size of the U_PDU or part that is included in this S_UNIDATA_REQUEST_CONFIRM Primitive.

The *U_PDU* argument in the S_UNIDATA_REQUEST_CONFIRM Primitive **shall** [6] be a copy of the whole or a fragment of the *U_PDU* argument of the S_UNIDATA_REQUEST Primitive for which the S_UNIDATA_REQUEST_CONFIRM Primitive is the response.

Using these arguments, the client will usually be able to uniquely identify the U_PDU that is being acknowledged. Depending on the implementation of the protocol, the last argument, *U_PDU,* may not be a complete copy of the original U_PDU but only a partial copy, i.e., only the first X bytes are copied for some value of *X*. If a partial U_PDU is returned, *U_PDU_response_frag_size* bytes **shall** [9] be returned to the client starting with the first byte of the U_PDU so that the client will have the U_PDU segment information. The number of bytes returned, *U_PDU_response_frag_size*, **shall**[10] be a configurable parameter in the implementation..

A.2.2.8   S_UNIDATA_REQUEST_REJECTED Primitive

Name :
S_UNIDATA_REQUEST_REJECTED
Arguments :
1. Reason
2. Destination Node Address
3. Destination SAP ID
4. Size of Rejected U_PDU (or part)
5. U_PDU (User Protocol Data Unit or part of it)
Direction :
Subnetwork Interface->Client

Description :
The S_UNIDATA_REQUEST_REJECTED primitive **shall** [1] be issued by the Subnetwork Interface Sublayer to inform a client that a S_UNIDATA_REQUEST was not delivered successfully.

This primitive **shall** [2] be issued if the client has requested Data Delivery Confirmation (either during Binding or for this particular U_PDU) and the data was unsuccessfully delivered. This primitive also **shall** [3] be issued to a client if a U_PDU larger than the MTU is submitted.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The argument *Reason* **shall** [(4)] specify why the delivery failed, using the encoding given in the table below:

Reason	Value
TTL Expired	1
Destination SAP ID not bound	2
Destination node not responding	3
U_PDU larger than MTU	4
Tx Mode not specified	5
Broadcast Not Allowed	6
Address Not Known	7
No ALE Mapping for Address	8
Soft Link Terminated for higher priority link	9
Soft Link Terminated to share channel	10
Idle Soft Link Terminated	11
Other	12

The value assigned to each reason **shall** be used to represent the reason in SIS Access Protocol (Annex S).

The *Destination Node Address* argument in the S_UNIDATA_REQUEST_REJECTED Primitive **shall** have the same meaning and be equal in value to the *Destination Node Address* argument of the S_UNIDATA_REQUEST Primitive for which the S_UNIDATA_REQUEST_REJECTED Primitive is the response.

The *Destination SAP_ID* argument in the S_UNIDATA_REQUEST_REJECTED Primitive **shall** [(7)] have the same meaning and be equal in value to the *Destination SAP_ID* argument of the S_UNIDATA_REQUEST Primitive for which the S_UNIDATA_REQUEST_REJECTED Primitive is the response.

The *Size of Rejected U_PDU* argument **shall** [(8)] be the size of the U_PDU or part that is included in this S_UNIDATA_REQUEST_REJECTED Primitive.

Just as specified for the S_UNIDATA_REQUEST_CONFIRM primitive, the *U_PDU* argument in the S_UNIDATA_REQUEST_REJECTED primitive may only be a partial copy of the original U_PDU, depending on the implementation of the protocol. If a partial U_PDU is returned, *U_PDU_response_frag_size* bytes **shall** [(9)] be returned to the client starting with the first byte of the U_PDU so that the client will have the U_PDU segment information. The number of bytes returned, *U_PDU_response_frag_size*, **shall** [(10)] be a configurable parameter in the implementation.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

A.2.2.9   S_UNIDATA_INDICATION Primitive

Name :
S_UNIDATA_INDICATION
Arguments :
1. Priority
2. Destination SAP ID
3. Destination Node Address
4. Transmission Mode
5. Source SAP ID
6. Source Node Address
7. Size of U_PDU
8. Number of Blocks in Error
9. Array of Block-Error Pointers
10. Number of Non-Received Blocks
11. Array of Non-Received-Block Pointers
12. U_PDU

Direction :
Subnetwork Interface->client
Description :
The S_UNIDATA_INDICATION primitive **shall** [1] be used by the Subnetwork Interface Sublayer to deliver a received U_PDU to the client.

The *Priority* argument **shall** [2] be the priority of the U_PDU.

The *Destination SAP ID* argument **shall** [3] be the SAP ID of the client to which this primitive is delivered.

The *Destination Node Address* argument **shall** [4] be the address assigned by the sending node to the U_PDU contained within this primitive. This normally will be the address of the local (i.e., receiving) node. It may however be a "group" address to which the local node has subscribed (Group Addresses and their subscribers are defined during configuration) and to which the source node addressed the U_PDU.

The *Transmission Mode* argument **shall** [5] be the mode by which the U_PDU was transmitted by the remote node and received by the local node; ie, ARQ, Non-ARQ
(Broadcast) transmission, or Non-ARQ w/ Errors.

The *Source SAP ID* **shall** [6] be SAP ID of the client that sent the U_PDU.

The *Source Node Address* **shall** [7] represent the node address of the client that sent the U_PDU.

The *Size of U_PDU* argument **shall** [8] be the size of the U_PDU that was sent and delivered in this S_UNIDATA_INDICATION S_Primitive.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The following four arguments **shall** [9] be present in the S_UNIDATA_INDICATION S_Primitive if and only if the Transmission Mode for the U_PDU is equal to Non-ARQ w/ Errors:

a) The *Number of Blocks in Error* argument **shall** [10] equal the number of data blocks in the U_PDU that were received in error by the lower layers of the subnetwork and that were passed on to the Subnetwork Interface Sublayer. This argument **shall** [11] specify the number of ordered pairs in the *Array of Block-Error Pointers* argument.

b) The *Array of Block-Error Pointers* argument **shall** [12] consist of a an array of ordered pairs, the first element in the pair equal to the location within the U_PDU of the data block with errors, and the second element equal to the size of the data block with errors.

c) The *Number of Non-Received Blocks* argument **shall** [13] equal the number of data blocks missing from the U_PDU because they were not received. This argument **shall** [14] specify the number of ordered pairs in the *Array of Non-Received-Block Pointers* argument.

d) The *Array of Non-Received-Block Pointers* **shall** [15] consist of an array of ordered pairs, the first element in the pair equal to the location of the missing data block in the U_PDU and the second element equal to the size of the missing data block.

The final argument, *U_PDU*, **shall** [16] contain the actual received user data for delivery to the client.

A.3   Peer-to-Peer Communication Protocols and S_PDUs

Peer Subnetwork Interface Sublayers, generally in different nodes, **shall** [1] communicate with each other by the exchange of Subnetwork Interface Sublayer Protocol Data Units (S_PDUs).

For the Subnetwork configurations currently defined in STANAG 5066, Peer-to-Peer Communication **shall** be [2] required for the Exchange of Client Data, including confirmation.

Explicit Peer-to-Peer communication **shall** [3] not be required for the establishment or termination of Soft Link or Broadcast Data Exchange sessions.

The Peer-to-Peer communication required for the exchange of Client Data is similar for all Data exchange sessions, using the facilities of lower sublayers in the protocol profile. The encoding of the S_PDUs and the protocol governing the Peer-to-Peer Communication are described in the following sections.

A.3.1   Subnetwork Interface Sublayer Protocol Data Units (S_PDUS) and Encoding Requirements

There are currently three types of S_PDUs. Additional S_PDU types may be defined in the future. The generic encoding of the eight S_PDU types showing the fields and subfields of the S_PDUs is shown in Figure A-1.
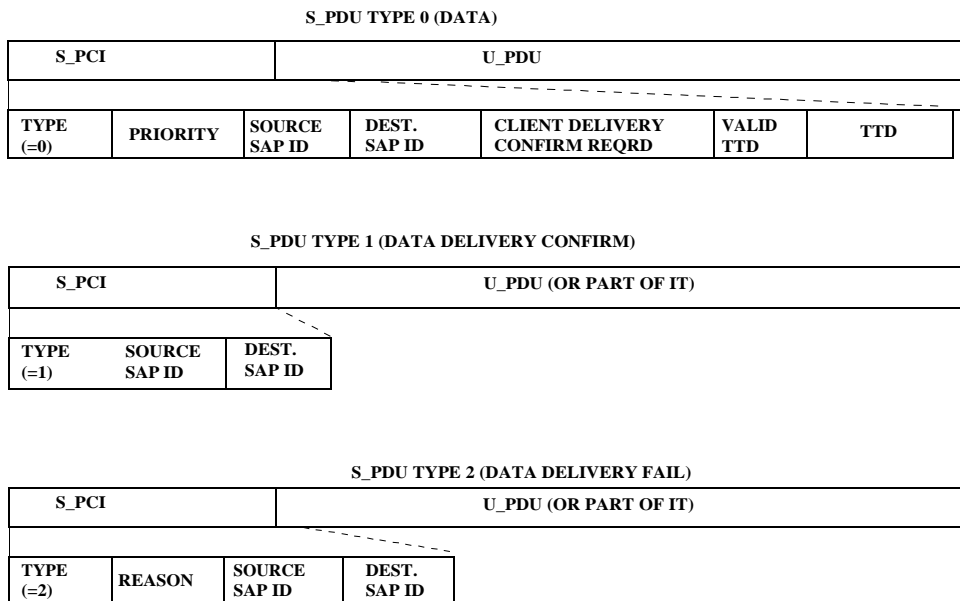
NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

S_PDU TYPE 0 (DATA)

S_PCI	U_PDU

TYPE (=0)	PRIORITY	SOURCE SAP ID	DEST. SAP ID	CLIENT DELIVERY CONFIRM REQRD	VALID TTD	TTD

S_PDU TYPE 1 (DATA DELIVERY CONFIRM)

S_PCI	U_PDU (OR PART OF IT)

TYPE (=1)	SOURCE SAP ID	DEST. SAP ID

S_PDU TYPE 2 (DATA DELIVERY FAIL)

S_PCI	U_PDU (OR PART OF IT)

TYPE (=2)	REASON	SOURCE SAP ID	DEST. SAP ID

Figure A-1: Generic Encoding of S_PDUs

The first encoded field **shall** [(1)] be common to all S_PDUs. It is called "TYPE" and **shall** [(2)] encode the type value of the S_PDU as follows:

S_PDU TYPE field value	S_PDU Name
0	DATA
1	DATA DELIVERY CONFIRM
2	DATA DELIVERY FAIL

The meaning and encoding of the remaining fields, if any, in an S_PDU **shall** [(3)] be as specified in the subsection below corresponding to the S_PDU type.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

A.3.1.1   DATA S_PDU

Type :
          "0" = DATA S_PDU
Encoding :

S_PCI	U_PDU

TYPE (=0)	PRIORITY	SOURCE SAP ID	DEST. SAP ID	CLIENT DELIVERY CONFIRM REQRD	VALID TTD	TTD

Figure A-2: Generic Encoding of the DATA S_PDU

	MSB 7	6	5	4	3	2	1	LSB 0
0	Type = 0				Priority			
1	Source SAP ID				Destination SAP ID			
2	CDCR	TTD Valid = 1	Not Used		MSB			
3	TTD							
4								LSB
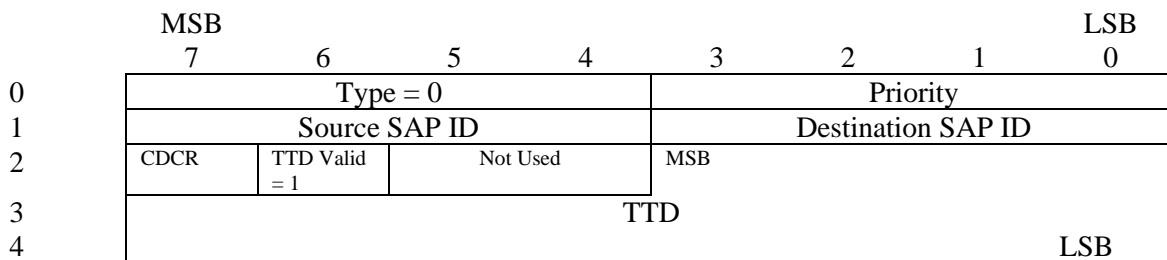
Figure A-3: Bit-Field Map of the DATA S_PDU S_PCI with valid TTD

	MSB 7	6	5	4	3	2	1	LSB 0
0	Type = 0				Priority			
1	Source SAP ID				Destination SAP ID			
2	CDCR	TTD Valid = 0	Not Used					

Figure A-4: Bit-Field Map of the DATA S_PDU S_PCI without TTD

Description :
          The DATA S_PDU **shall** [1] be transmitted by the Subnetwork Interface Sublayer in order to send client data to a remote peer sublayer.

          The DATA S_PDU shall (2) be encoded as specified in Figure A-2, Figure A-3, Figure A-4 and in the paragraphs below.

          This S_PDU **shall** [3] consist of two parts:
                    a)   the first part **shall** [4] be the S_PCI (Subnetwork Interface Sublayer Protocol Control Information) and represents the overhead added by the sublayer;
                    b)   the second part **shall** [5] be the actual client data (U_PDU).

          The first field of four bits the S_PCI part **shall** [6] be "TYPE". Its value **shall** [7] be equal to 0 and identifies the S_PDU as being of type DATA.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The second field of four bits **shall** [8] be "PRIORITY" and represents the priority of the client's U_PDU. The "PRIORITY" field **shall** [9] be equal in value to the corresponding argument of the S_UNIDATA_REQUEST primitive submitted by the client..

The third field of four bits of the S_PCI **shall** [10] be the "SOURCE SAP ID" and identifies the client of the transmitting peer which sent the data.

The fourth field of four bits **shall** [11] be the "DESTINATION SAP ID" and identifies the client of the receiving peer which must take delivery of the data. There is no need to encode the source and destination node addresses in the S_PDU as this information is relayed between the peers by the underlying sublayers. The "DESTINATION SAP ID" **shall** [12] be equal in value to the corresponding argument of the S_UNIDATA_REQUEST primitive submitted by the client

The fifth field of the S_PCI **shall** [13] be "CLIENT DELIVERY CONFIRM REQUIRED", and is encoded as a single bit that can take the values "YES" (=1) or "NO" (=0). The value of this bit **shall** [1 be set according to the *Delivery Mode* requested explicitly for this U_PDU (see S_UNIDATA_REQUEST Primitive), which **may** be defaulted to the *Service Type* requested by the sending client during binding (see S_BIND_REQUEST primitive).

The sixth field **shall** [15] be the VALID TTD field, and is encoded as a single bit that can take the values "YES" (=1) or "NO" (=0), indicating the presence of a valid TTD within the S_PCI.

The seven field of the S_PCI **shall** [14] be two unused bits that are reserved for future use.

The eighth and last field of the S_PCI **shall** [15] be "TTD" and represents the TimeToDie for this U_PDU. The first four bits of this field **shall** [16] have meaning if and only if the VALID TTD field equals "YES", the remaining 16 bits of the field **shall** [17] be present in the S_PCI if and only if the VALID TTD field equals "YES".

The TTD field encodes the Julian date modulo 16, and the GMT in seconds after which time the S_PDU must be discarded if it has not yet been delivered to the client. The simple Julian date system, which numbers the days of the year consecutively starting with 001 on 1 January and ending with 365 on 31 December (or 366 on leap years). The Julian date modulo 16 part of the TTD **shall** be mapped into the first four bits of the TTD field (i.e., bits 0-3 of byte 2 of the S_PDU).

The 16 high bits of the GMT part of the TTD shall be mapped into the 2 remaining bytes of the TTD field; the LSB of the GMT shall be discarded. If the "VALID TTD" flag bit of a DATA S_PDU is set (=1) then the complete TTD 20-bit field is present and its value must be used. If this flag bit is not set (=0), the last two bytes of the TTD field are not present (to conserve overhead) and the TTD must not be used. The "VALID TTD" flag bit allows the transmitting peer to specify whether the receiving peer should discard the S_PDU by based on TTD or it delivered the U_PDU to the client without consideration of the TTD.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

A.3.1.2   DATA DELIVERY CONFIRM S_PDU

Type :

     "1" = DATA DELIVERY CONFIRM

Encoding :

S_PDU TYPE 1 (DATA DELIVERY CONFIRM)

S_PCI	U_PDU (OR PART OF IT)

TYPE (=1)	SOURCE SAP ID	DEST. SAP ID

Figure A-5: Generic Encoding of the DATA DELIVERY CONFIRM S_PDU

|	MSB							LSB
	7	6	5	4	3	2	1	0
0	Type = 1				Priority			
1	Source SAP ID				Destination SAP ID			

Figure A-6: Bit-Field Map of the DATA DELIVERY CONFIRM S_PDU S_PCI

Description :

     The DATA DELIVERY CONFIRM S_PDU **shall** be [(1)] transmitted in response to a successful delivery to a Client of a U_PDU which was received in a DATA type S_PDU in which the "CLIENT DELIVERY CONFIRM REQUIRED" field was set to "YES". The DATA DELIVERY CONFIRM S_PDU **shall** be [(2)] transmitted by the Subnetwork Interface Sublayer to the peer sublayer which originated the DATA type S_PDU.

     The first part of the DATA DELIVERY CONFIRM S_PDU **shall** [(3)] be the S_PCI, while the second part **shall** [(4)] be a full or partial copy of the U_PDU that was received and delivered to the destination Client.

     The first field of the S_PCI part **shall** [(5)] be "TYPE" and its value **shall** [(6)] equal 1 to identify the S_PDU as being of type DATA DELIVERY CONFIRM.

     The remaining fields and their values for the S_PCI part of the DATA DELIVERY CONFIRM S_PDU **shall** [(7)] be equal in value to the corresponding fields of the DATA S_PDU for which this DATA DELIVERY CONFIRM S_PDU is a response.

     The peer sublayer that receives the DATA DELIVERY CONFIRM **shall** [(8)] inform the client which originated the U_PDU that its data has been successfully delivered to its Destination by issuing a S_UNIDATA_REQUEST_CONFIRM in accordance with the data exchange protocol of Section A.3.2.3.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

A.3.1.3   DATA DELIVERY FAIL S_PDU

Type :
          "2" = DATA DELIVERY FAIL
Encoding :

S_PDU TYPE 2 (DATA DELIVERY FAIL)

S_PCI	U_PDU (OR PART OF IT)

TYPE (=2)	REASON	SOURCE SAP ID	DEST. SAP ID

Figure A-7: Generic Encoding of the DATA DELIVERY FAIL S_PDU

	MSB							LSB
	7	6	5	4	3	2	1	0
0	Type = 2				Priority			
1	Source SAP ID				Destination SAP ID			

Figure A-8: Bit-Field Map of the DATA DELIVERY FAIL S_PDU S_PCI

Description :
          The DATA DELIVERY FAIL S_PDU **shall** [1] be transmitted in response to a failed delivery to a Client of a U_PDU that was received in a DATA type S_PDU with the "CLIENT DELIVERY CONFIRM REQUIRED" field set to "YES".

          The first part of this S_PDU **shall** [2] be the S_PCI.

          The second part **shall** [3] be a full or partial copy of the U_PDU that was received but not delivered to the destination Client.

          The first field of the S_PCI **shall** [4] be "TYPE". Its value **shall** [5] be equal to 2 and identifies the S_PDU as being of type DATA DELIVERY FAIL.

          The second field **shall** [6] be "REASON" and explains why the U_PDU failed to be delivered. It can take a value in the range 0-15; valid reasons are defined in the table below.

Reason	Value
Unassigned and reserved	0
Destination SAP ID not bound	1
Unassigned and reserved	2-15

          The SOURCE SAP_ID and DESTINATION SAP_ID fields of the S_PCI **shall** [7] be equal in value to the corresponding fields of the DATA S_PDU for which the DATA DELIVERY FAIL S_PDU is a response.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The peer sublayer that receives the DATA DELIVERY FAIL S_PDU, **shall** [8] inform the client which originated the U_PDU that its data was not delivered to the destination by issuing a S_UNIDATA_REQUEST_REJECTED primitive, in accordance with the data exchange protocol of Section A.3.2.3.

A.3.2   Peer-to-Peer Communication Protocol

This section specifies the protocols governing the Peer-to-Peer communication for Establishing and Terminating Soft Link Data Exchange Sessions, Establishing and Terminating Broadcast Data Exchange Sessions and Exchanging Client Data. In these specifications, the node whose local client or Subnetwork Interface Sublayer first requests a Data Exchange Session is denoted as the caller or calling node and the remote node is denoted as the called node.

A.3.2.1   Soft Link Data Exchange Session

The Subnetwork Interface Sublayer initiates Soft Link Data Exchange Sessions with remote peers based on the destinations of queued client U_PDUs. In particular, sublayer management algorithms must be established to initiate the protocols for establishment or termination a Soft Link Data Exchange Session. This STANAG allows these sublayer management algorithms to be based on implementation dependent criteria and factors. The use of comparative U_PDU queue-length for given clients, source-destination sets and priority levels for any implementation is allowed and expected (even if the algorithms are trivial) but remain beyond the scope of this STANAG.

A.3.2.1.1   Soft Links for Non-ARQ data

Soft links are always established prior to transfer of ARQ data over the CAS.

When CAS is used in Multiple Simultaneous Peer Access, as described in Annex B, Non-ARQ will always be transferred without a Softlink.

When CAS is used in Single Peer Access Mode with 1:1 ALE, a soft link **shall** be established for all data.   The CAS signals this to the SIS by rejecting Non-ARQ data with reason "Physical Link Needed". SIS layer **shall** establish as soft link and then resend the non-ARQ data.

A.3.2.1.2   Protocol for Establishing a Soft Link Data Exchange Session

The establishment of Soft Link Data Exchange Sessions **shall** [1] not require explicit peer-to-peer handshaking within the Subnetwork Interface Sublayer.

The calling peer **shall** [2] implicitly establish a Soft Link Data Exchange Session by requesting its Channel Access Sublayer to make a physical link to the required remote node, using the procedure for making physical links specified in Annex B. In accordance with these procedures, both peer Subnetwork Interface Sublayers (i.e., the calling and called sublayers) are informed about the successful making of a physical link between their nodes by their respective Channel Access Sublayers.

After the physical link is made, both peer Subnetwork Interface Sublayers **shall** [3] declare that the Soft Link Data Exchange Session has been established between the respective source and destination nodes. Data may then be exchanged in accordance with the protocols specified in Section A.3.2.4.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

A.3.2.1.3   Protocol for Terminating a Soft Link Data Exchange Session

No peer-to-peer communication by the Subnetwork Interface Sublayer **shall** [(1)] be required to terminate a Soft Link Data Exchange Session.

A Soft Link Data Exchange Session **shall** [(2)] be terminated by either of the two peers by a request to its respective Channel Access Sublayer to break the Physical Link in accordance with the procedure specified in Annex B. Both Subnetwork Interface sublayers will be informed about the breaking of the Physical link by their respective Channel Access Sublayers.

Since a called peer can terminate a Soft Link Data Exchange Session if it has higher priority data destined for a different Node, called peers **shall** [(3)] wait a configurable minimum time before unilaterally terminating sessions, to prevent unstable operation of the protocol.

> Note: The caller sublayer normally initiates the termination of the session (by breaking the physical link) based on the destinations of its queued U_PDUs, and on any ongoing communication with the distant node. The inter-layer signaling for coordination would normally be carried out via the subnetwork management sublayer. The called sublayer can also terminate the session if it has high priority data destined for a different node. However, called sublayers should wait a configurable minimum time before unilaterally terminating sessions, otherwise an unstable condition may arise if all nodes in the network have data to transmit and called sublayers immediately close sessions in order to establish other sessions as callers. If such a situation arises, the efficiency of a subnetwork will deteriorate as a result of nodes continuously establishing and terminating sessions without actually transmitting data. The minimum amount of time that a called sublayer should wait before it attempts to terminate a Soft Link Session must be carefully chosen and will depend on a number of factors such as the subnetwork size and configuration. Specification of this and other parameters as a configurable but required value allows implementations of the STANAG to be tuned for specific network, with the values for these parameters distributed as part of the standard operating procedures for a given network.

After the Subnetwork Interface Sublayer has been notified that the Physical Link has been broken, the Subnetwork Interface Sublayer **shall** [(4)] declare the Soft Link Exchange Session as terminated.

A.3.2.1.4   Queue Management

The intent of this STANAG, as noted above, is to maximize flexibility of implementation when managing queues.

When the CAS follows the Multiple Simultaneous Peer Access model described in Annex B, the SIS layer will be able to establish multiple queues to different peers, each with an open soft link.   In this model the DTS level queues will handle D_PDUs of different priority according to priority.   Once a soft link is open, the SIS **may** pass all data to the DTS and allow DTS to handle priority and sharing of resource to transmit data to different peers.   Annex C, section 7.5.2 describes DTS handling of priority and peer sharing.

When the Single Peer Access model is used in a deployment with multiple peers it will be necessary to perform queueing at the SIS level.   This is because a limited number of soft links can be open at a time, and the SIS will need to queue data for peers where the soft link is not open.   The SIS will need to control which links are open. The following rules apply:

- If a soft link is open and traffic of a higher priority than any being transmitted arrives, the SIS **shall** close active link as quickly as possible, noting the considerations in Section A.3.2.1.3.
- Where there is traffic for multiple peers, the SIS **shall** ensure fairness between the peers and open soft links to each peer in turn at reasonable intervals.   Traffic to one peer **shall not** be allowed to prevent traffic to other peers.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

The SIS **shall** close a soft link using the CAS C_PHYSICAL_LINK_BREAK primitive with reason "Soft Link Terminated for higher priority link" or "Soft Link Terminated to share channel" or "Idle Soft Link Terminated".

Data rejected by the DTS will return D_UNIDATA_REQUEST_REJECTED to the CAS, which will be passed to the SIS as C_UNIDATA_REQUEST_REJECTED. When this is due to soft link termination, the reason for the reject **shall** match the reason given by SIS for closing the soft link. The SIS will need to make a balance between rapid switching of soft link and allowing transfers on soft link to complete before closing the soft link. This is an implementation choice.

A.3.2.2   Protocol for Establishing and Terminating a Broadcast Data Exchange Session

No explicit peer-to-peer communication **shall** [1] be required to establish and terminate a Broadcast Data Exchange Session. A Broadcast Data Exchange Session is established and terminated either by a management

process or unilaterally by the Subnetwork Interface Sublayer based on a number of criteria as explained in section A.1.1.3.

As noted in section A.1.1, clients may interleave requests for data-exchange sessions. At some point, the subnetwork might also be configured to provide exclusive support for a Broadcast Data Exchange Session. In this case, when the subnetwork is first configured by the local (implementation-dependent) management function to provide exclusive support for a Broadcast Data Exchange Session the Subnetwork Interface Sublayer **shall** [2] send an S_UNBIND_INDICATION to any bound clients that had requested ARQ Delivery Service, with the REASON = "ARQ Mode Unsupportable during Broadcast Session". Subsequent S_BIND requests by clients requesting ARQ service **shall** [3] be rejected with the same reason.

A.3.2.3   Protocol for Exchanging Client Data

After a Data Exchange Session of any type has been established, sublayers with client data to exchange **shall** [1] exchange DATA (TYPE 0) S_PDUs using the protocol specified below and in accordance with the service characteristics of the respective session.

The sublayer **shall** [2] discard any U_PDU submitted by a client where the U_PDU is greater in size than the Maximum Transmission Unit (MTU) size assigned to the client by the S_BIND_ACCEPTED Primitive issued during the client-bind protocol.

If a U_PDU is discarded because it exceeded the MTU size limit and if the DELIVERY CONFIRMATION field for the U_PDU specifies CLIENT DELIVERY CONFIRM or NODE DELIVERY CONFIRM, the sublayer **shall** [3] notify the client that submitted the U_PDU as follows:
- the sublayer **shall** send a S_UNIDATA_REQUEST_REJECTED Primitive to the client;
- the REASON field **shall** be equal to "U_PDU Larger than MTU".

For U_PDUs that have been accepted for transmission, the sending sublayer retrieves client U_PDUs and their associated implementation-dependent service attributes (such as the S_Primitive that encapsulated the U_PDU) from its queues (according to Priority and other implementation-dependent criteria), and proceeds as follows:

- the sending sublayer **shall** [7] encode the retrieved U_PDU into a DATA (TYPE 0) S_PDU, transferring any service attributes associated with U_PDU to the S_PDU as required;

- the sending sublayer **shall** [8] encode the resulting DATA (TYPE 0) S_PDU in accordance with the C_Primitive interface requirements of the Channel Access Sublayer as specified in Annex B, i.e,:

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

- the sublayer **shall** [9] encode the S_PDU as a C_UNIDATA_REQUEST Primitive of the priority corresponding to that initially specified by the client in the S_Primitive, otherwise;

- the sending sublayer then **shall** [11] pass the resulting C_primitive to the Channel Access Sublayer for further processing to send the DATA (TYPE 0) S_PDU to its remote peer.

- if the service attributes for the U_PDU require NODE DELIVERY CONFIRMATION, the sublayer **shall** [12] wait for a configurable time for a response as follows:

- if the sublayer receives a C_UNIDATA_REQUEST_CONFIRM prior to the end of the waiting time, the sublayer **shall** [13] send to the client a S_UNIDATA_REQUEST_CONFIRM Primitive;

- otherwise, if the sublayer receives a C_UNIDATA_REQUEST_REJECTED the sublayer **shall** [14] send to the client a S_UNIDATA_REQUEST_REJECTED Primitive, unless the reject reason is "Physical Link Needed". If the rejection is for this reason (which is expected only for Non-ARQ data), the sublayer **shall** establish a soft link and resubmit the C_UNIDATA_REQUEST;

- otherwise, if the waiting time ends prior to receipt of any response indication from the Channel Access sublayer, the Subnetwork Interface sublayer **shall** [15]
 send to the client a S_UNIDATA_REQUEST_REJECTED Primitive.
The REASON field shall be set equal to "Destination Node Not Responding".

- if the service attributes for the U_PDU require CLIENT DELIVERY CONFIRMATION, the sending sublayer shall (16) wait for a configurable time for a response as follows:

- if the Subnetwork Interface sublayer receives a C_Primitive confirming node-node delivery (i.e., a C_UNIDATA_REQUEST_CONFIRM Primitive) and a "DATA DELIVERY CONFIRM" (TYPE 1) S_PDU is received from the remote sublayer prior to the end of the waiting time, the Subnetwork Interface sublayer **shall** [17] send to the client a S_UNIDATA_REQUEST_CONFIRM Primitive;

- otherwise, if the Subnetwork Interface sublayer receives either a "reject" C_Primitive from the Channel Access Sublayer or a "DATA DELIVERY FAIL" (TYPE 2) S_PDU from the remote peer prior to the end of the waiting time, the Subnetwork Interface sublayer **shall** [18] send to the client either a S_UNIDATA_REQUEST_REJECTED Primitive. Tthe REASON field **shall** [18.1] be derived from the "DATA DELIVERY FAIL" (TYPE 2) S_PDU or the reject C_Primitive that was received;

- otherwise, if the waiting time ends prior to receipt of a response message, the sublayer **shall** [19] send to the client a S_UNIDATA_REQUEST_REJECTED Primitive.,The REASON field shall be set equal to "Destination Node Not Responding".

- On completion of these actions by the sending sublayer the client data delivery protocol terminates for the given DATA (TYPE 0) S_PDU.

A receiving sublayer manages the client data exchange protocol as follows:

- the receiving sublayer **shall** [20] accept encoded DATA (TYPE 0) S_PDUs from the Channel Access Sublayer using C_Primitives in accordance with the interface requirements specified

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

in Annex B.

- the receiving sublayer **shall** [21] extract the U_PDU, Destination SAP_ID and the other associated service attributes from the DATA (TYPE 0) S_PDUs as required;

- if there is no client bound to the destination SAP_ID, the receiving sublayer **shall** [22] discard the U_PDU by; otherwise,

- the sublayer shall (23) deliver the extracted U_PDU to the destination client bound to Destination SAP_ID using a S_UNIDATA_INDICATION Primitive;

- if the received S_PDU has the "CLIENT DELIVERY CONFIRM REQUIRED" field set equal to "YES", then the sublayer **shall** [25] provide delivery confirmation as follows:

    - if a client was bound to the Destination SAP_ID, the sublayer **shall** [26] encode as required and send a "DATA DELIVERY CONFIRM" (TYPE 1) S_PDU to the sending sublayer;

    - if a client was not bound to the Destination SAP_ID, the sublayer **shall** [27] encode as required and send a "DATA DELIVERY FAIL" (TYPE 2) S_PDU to the sending sublayer.

- On completion of these actions by the receiving sublayer the client data delivery protocol terminates for the given DATA (TYPE 0) S_PDU.

Implementation-dependent queuing disciplines, flow-control procedures, or other characteristics in the sublayer **shall** [28] not preclude the possibility of managing the data exchange protocol for more than one U_PDU at a time. In particular, the Subnetwork Interface Sublayer **shall** [29] be capable of sending a U_PDU, encapsulated in a DATA (TYPE 0) S_PDU and C_Primitive as required, to the Channel Access Sublayer prior to receipt of the data-delivery-confirm response for a U_PDU sent earlier.

> [Note: This requirement mitigates the reduction in link throughput that occurs when a subnetwork ceases transmission of any U_PDUs while it awaits confirmation of their delivery. The performance degradation is typical of that which occurs when using a STOP-AND-WAIT form of ARQ protocol anywhere in a communication system.]

The nominal procedures for exchanging DATA S_PDUs for both the Sending and Receiving Peers are shown in Figure A-9 and Figure A-10. This STANAG acknowledges that other implementations may satisfy the requirements stated above.
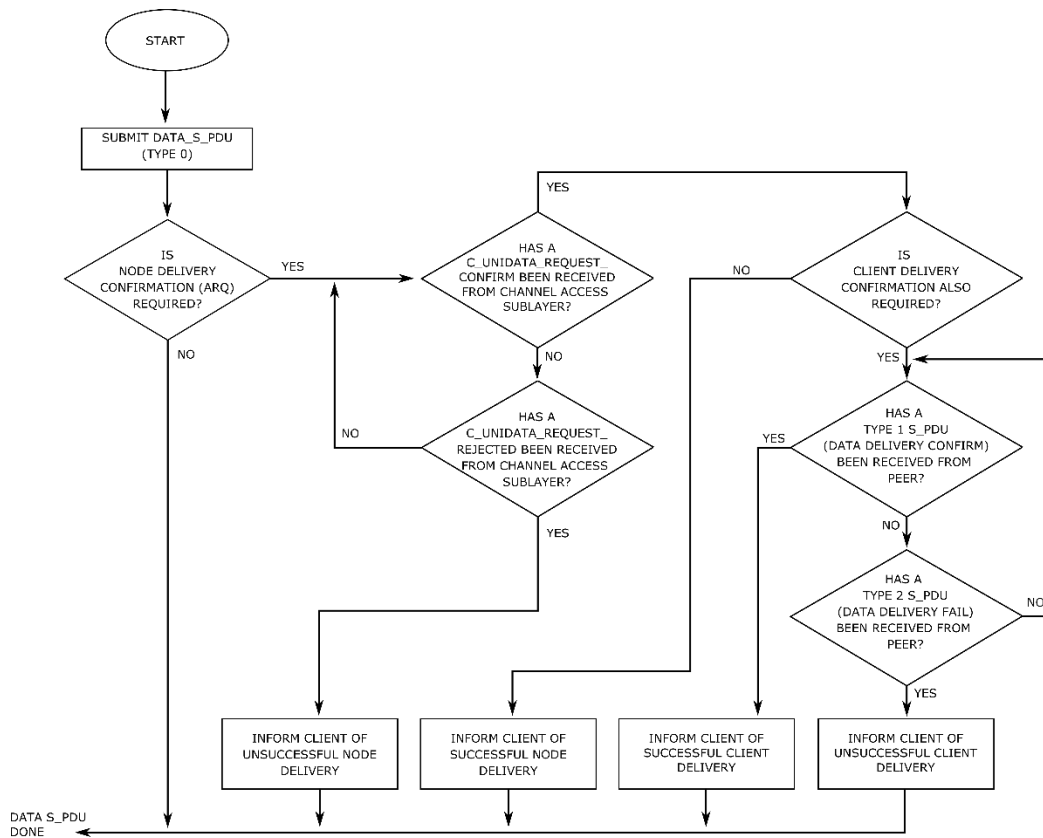
NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

Figure A-9: Data Exchange Procedures: SENDING PEER

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW
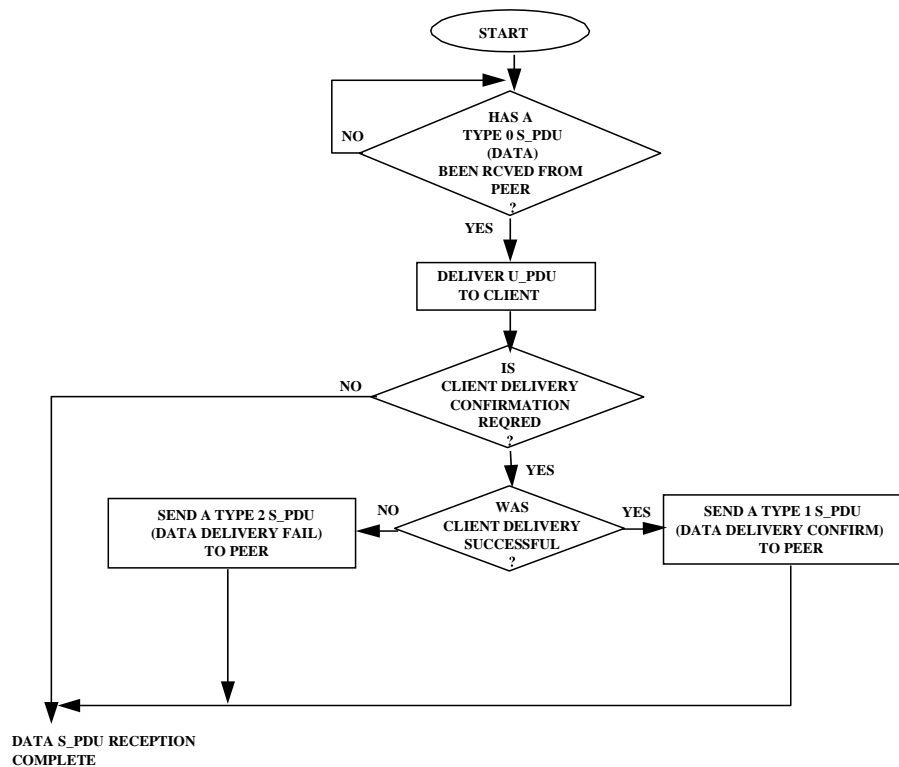
ANNEX A
STANAG 5066
Edition 4

Figure A-10: Data Exchange Procedures: RECEIVING PEER

A.4   Changes in Edition 4

The following changes are made relative to Edition 3.

1.   The SIS protocol is moved into it's own Annex (Annex S) rather than being split between Annex's A and F. This improves clarity of Annex A and allows it to follow OSI model of service specification with peer protocol.

2.   Hard Links are removed. These added significant complexity/confusion and are of no current or anticipated use.

3.   Expedited data is removed. This is of no current or anticipated operational use. Expedited data was used in some OSI protocols, but is not used in modern applications. Edition 3 Annex H also recommends against its use.

4.   Rank is removed. This apparently useful capability is confusing, as its use (excluding hard links) is only for two functions that are not needed:

   a.   Restriction of Management Messages to Rank 15. This is a useless and strange control, as rank is client selected.
   b.   Choice of client to use if resources are constrained (from Annex H). This is not relevant to a modern implementation.

5.   Addition of new error types.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX A
STANAG 5066
Edition 4

6.   Clarification of queueing model to support ALE use as defined in Annex B.

7.   Prevent use of "in-order" with Non-ARQ.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

Annex B: Channel Access Sublayer
(mandatory)
DRAFT 1.5 of 9th April 2020

The channel access sublayer handles links between peers.

B.1. <u>Channel Access Sublayer Service Definition</u>

The Channel Access Sublayer provides services to the Subnetwork Interface Sublayer. These services are:

B.1.1.1. Execute requests by the Subnetwork Interface sublayer to *"Make"* and *"Break"* Physical Links

B.1.1.2. Notify the Subnetwork Interface sublayer of changes in the state of a Physical Link.

B.1.1.3. Accept S_PDUs (encapsulated in the appropriate primitive) from the Subnetwork Interface sublayer for transmission on a Physical Link.

B.1.1.4. Deliver S_PDUs (encapsulated in the appropriate primitive) received on a Physical Link to the Subnetwork Interface sublayer.

In order to provide these services, the Channel Access Sublayer implements a protocol that specifies the tasks that must be executed and the rules that must be obeyed by the sublayer. While a number of different channel-access protocols are possible, the one that is suitable for this document is referred to as the Channel Access Type 1 (CAS-1) protocol, and described herein.

The Channel Access Sublayer also manages control of 1:1 ALE links between peer nodes.

B.2. <u>Interface Primitives Exchanged with the Subnetwork Interface Sublayer</u>

The implementation of the interface between the Channel Access Sublayer and the Subnetwork Interface sublayer is not mandated or specified by this STANAG. Since the interface is internal to the subnetwork architecture and may be implemented in a number of ways it is considered beyond the scope of STANAG 5066. A model of the interface has been assumed, however, for the purposes of discussion and specification of other sublayer functions.

Despite the advisory nature of the conceptual model of the internal interface between the Subnetwork Interface sublayer and the Channel Access Sublayer, there are some mandatory requirements that are placed on any interface implementation.

The interface must support the service-definition for the Channel Access Sublayer, i.e.:

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

1. the interface **shall** enable the Subnetwork Interface sublayer to submit requests to change the state of a physical link, i.e., to make or break a physical link with a specified node address;

2. The interface **shall** allow the Channel Access sublayer to accept S_PDUs from the Subnetwork Interface sublayer

3. The interface **shall** allow the Channel Access sublayer to deliver S_PDUs to the Subnetwork Interface sublayer.

Additionally, the protocol-control information from the Subnetwork Interface sublayer that is required for the management of the Channel Access sublayer **shall** not be derived from knowledge of the contents or format of any client data or U_PDUs encapsulated within the S_PDUs exchanged over the interface.
[Note: user's that encrypt their traffic prior to submittal may use the subnetwork. Subnetwork operation must be possible with client data in arbitrary formats that are unknown to the subnetwork, therefore any service requirements or indications must be provided by interface control information provided explicitly with the user data.]

The Channel Access Sublayer does not queue traffic. The DTS will queue traffic and in some modes of operation SIS will queue traffic. CAS can process information with its peers and between the layers immediately, without need for any queueing.

The interface may use knowledge of the contents of S_PDUs (excluding the contents of any encapsulated U_PDUs) to derive protocol control information for the Channel Access sublayer. This approach is highly discouraged, however. The recommended approach for implementation is that information required for protocol control within the Channel Access sublayer should be provided explicitly in appropriate interface primitives.

In keeping with accepted practice in the definition of layered protocols, and as a means for specifying the operations of the sublayers that are mandated by this STANAG, the communication between the Channel Access Sublayer and the Subnetwork Interface sublayer is described herein with respect to a set of Primitives. The interface Primitives are a set of messages for communication and control of the interface and service requests made between the two layers.

By analogy to the design of the client-subnetwork interface, the technical specification of the Channel Access Sublayer assumes communication with the Subnetwork Interface sublayer using primitives prefixed with a "C_". A minimal set of C_Primitives has been assumed that meet the requirements stated above and the general function for each C_Primitive is given in Table B-1. These C_Primitives are given without benefit of a list of arguments or detailed description of their use. As noted initially, they are offered for information only as a means of describing the interaction between the Channel Access sublayer and the Subnetwork Interface sublayer, and as general guidance for a possible implementation.

Table B-1- Nominal Definition of C_Primitives for the Interface between the Channel Access sublayer and the Subnetwork Interface sublayer (non-mandatory, for information-only)

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

NAME OF PRIMITIVE	DIRECTION (see Note)	COMMENTS
C_PHYSICAL_LINK_MAKE	SIS →CAS	request to make a physical link of a specified type to a specified node
C_PHYSICAL_LINK_MADE	CAS→SIS	report that a physical link was made
C_PHYSICAL_LINK_REJECTED	CAS→SIS	report that a physical link could not be made and that the initiating request is rejected
C_PHYSICAL_LINK_BREAK	SIS→CAS	request to break a physical link
C_PHYSICAL_LINK_BROKEN	CAS→SIS	report that a physical has been broken (by request, or unilaterally)
C_UNIDATA_REQUEST	SIS→CAS	delivers an S_PDU to the Channel Access Sublayer, requesting normal data-delivery service
C_UNIDATA_REQUEST_CONFIRM	CAS→SIS	confirms S_PDU delivery to the remote node using the normal data-delivery service
C_UNIDATA_REQUEST_REJECTED	CAS→SIS	notifies that an S_PDU could not be delivered to a remote node using the normal data-delivery service
C_UNIDATA_INDICATION	CAS→SIS	delivers an S_PDU that had been received using the normal delivery service to the Subnetwork Interface sublayer

[Note: SIS = Subnetwork Interface Sublayer; CAS = Channel Access Sublayer;
→ = direction of message flow from source to destination sublayer]

B.2.1. Errors from C_ Primitives

The parameters of the C_ service primitives are clear from the definitions are not listed explicitly. This section lists errors from the service primitives.

B.2.1.1. Errors from C_PHYSICAL_LINK REJECTED

The errors are grouped by where they came from:

1. Rejection by peer with sub-reasons from PHYSICAL LINK REJECTED C_PDU:
   a. Reason Unknown
   b. Broadcast only node
2. DTS rejected transfer of Link Make using D_ EXPEDITED_UNIDATA_REQUEST_REJECTED. The DTS has no peer mechanism for rejection. The only local mechanism for rejection is TTL expiry. TTL **should** be set to value large enough to ensure this does not happen in practice.
3. Timer expired after final transmission of PHYSICAL LINK REQUEST C_PDU.
4. Local Reasons, defined by this specification:
   a. "Maximum Physical Links In Use".

These errors are handled as follows:

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

1. Broadcast only reflects a mis-configuration which needs to be addressed.

2. "Maximum Physical Links In Use" will lead to SIS layer holding off on establishing a physical link to this peer until another physical link has been closed.

3. Other errors will lead to the SIS layer rejecting all S_PDUs queued for the link, and allowing the SIS user to resubmit the messages if it wishes.

B.2.1.2. Errors from C_UNIDATA_REQUEST_REJECTED

C_UNIDATA_REQUEST is an unacknowledged service, and so there are no peer errors. The following errors are noted:

1. DTS rejected transfer D_ UNIDATA_REQUEST_REJECTED. The DTS has no peer mechanism for rejection. The only local mechanism for rejection of a valid PDU is:

    a. TTL expiry.

2. Local reasons, defined in this specification

    a. Broadcast Not Allowed

    b. Address Not Known

    c. No ALE Mapping for Address

    d. Physical Link Needed

TTL Expiry and the first three local reasons will be passed up to the SIS user along with rejection of the S_PDU. The "Physical Link Needed" error is used when Non-ARQ data is sent to a peer using ALE, where it is necessary to establish a physical link for all traffic to the peer.

B.3.   Models of Peer Access

There are two models of channel access, described in the next two sections:

1. Multiple Simultaneous Peer Access

2. Single Peer Access

The model in use **shall** be the same for all nodes using a channel.

For a profile of this standard, the model or models supported need to be specified.

B.4.   Multiple Simultaneous Peer Access

The first model is that a node on a channel can simultaneously transmit to multiple peers, utilizing the

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

broadcast nature of the underlying channel.   Control of access to the channel is then managed by the MAC layer using CSMA (Annex K) or Wireless Token Ring Protocol (Annex L).

In this model, the Channel Access Sublayer can submit data for all destinations directly to the DTS.

B.5.   Single Peer Access

In this second model, the Channel Access Sublayer accesses a channel or channels that can communicate with just one peer. This model **shall** be used in two situations.

1. Where Automatic Link Establishment (ALE) is used to link 1:1 between peers.   When used, this ALE 1:1 mode of operation **shall** be configured for all nodes accessing the channel; and/or

2. Duplex operation, where exchange of data between two nodes use two exclusive channels, one for transmission and one for reception.

In this model, the Channel Access Sublayer will control the active peer (or active peers with each peer on an independent channel).   Requests for transfers to other peers will be rejected, requiring the SIS to queue data and manage traffic for multiple peers where the number of channels available is less than the number of peers for which there is traffic.

B.5.1.   Duplex without ALE

When Duplex configuration is used without ALE, a pair of channels need to be set up between two nodes (one for each direction of transfer).   This setup is only viable for one pair of nodes.

For this configuration, Channel Access operates following the procedures of multiple simultaneous peer access.   No queuing a the CAS layer is needed, as there is only one peer.   One additional check is made in this configuration:

1. If data is addressed to an any unicast STANAG 5066 address other than that of the single peer, it **shall** be rejected with reason code "Address Not Known"

Group addressing **may** be used. Although it is not usually of operational benefit when there is a single peer, it may be helpful for some applications.

B.5.2.   Operation with 1:1 ALE

A channel may be configured to use ALE (Automatic Link Establishment) to connect to one peer at a time. The procedure defined here can be used with any ALE protocol, including 2G, 3G and 4G. When ALE is used, it **shall** be used for all communication.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

A node **may** be configured to have multiple channels of this nature. Where there are multiple channels, the DTS for each channel **shall** operate independently.

When data arrives for a peer (ARQ or non-ARQ), the following checks are made:

1. If the data is addressed to a broadcast address, it **shall** be rejected with reason code "Broadcast Not Allowed".
2. If data is addressed to an unknown STANAG 5066 address, it **shall** be rejected with reason code "Address Not Known"
3. If an ALE address cannot be determined for the address, it **shall** be rejected with reason code "No ALE Mapping for Address"

When an ALE connection is established to a peer, a CAS-1 link is implicitly established for the peer. The CAS-1 soft link establishment procedure described below **may** be followed, either to determine which edition of STANAG 5066 a peer supports or to negotiate the optional use of short frame sequence number D_PDUs between Edition 4 peers. As ALE connection initializes the CAS-1 link, the initialization procedure does not need to be used.

When closing an ALE link with a peer, the CAS-1 link break procedure **shall** be followed, even if the ALE link has only been used to carry non-ARQ traffic. This ensures that both local system and peer know that the CAS-1 link is finished and can either return to ALE searching or establish an ALE link to another peer.

If the ALE link is closed externally (e.g., by an operator) and the above procedure is not followed, the CAS-1 link **shall** be treated as closed.

The Channel Access Sublayer **shall** maintain a list of open physical links with maximum size of the number of channels. Each channel will have an independent DTS. When there are no spare channels, any request to open a new physical link **shall** be rejected with error "Maximum Physical Links In Use". Commonly this will be limited to one channel.

When ALE is in use, non-ARQ traffic will need to use ALE as well. To achieve this, a physical links is mandated for non-ARQ traffic as well as ARQ traffic. In the event of non-ARQ traffic being submitted when there is not a physical link to the destination, the submission **shall** be rejected with reason "Physical Link Needed".

B.5.3. Operation with Multicast ALE

It is also possible to use ALE to open a link to multiple nodes. This **may** be done when multicast traffic is processed, for example in support of a multicast protocol such as ACP 142. When an ALE link is opened to multiple nodes a CAS-1 link is established with each peer node, which **may** use explicit exchange of protocol.

The CAS-1 links for each node **shall** be terminated together, so that the ALE link can be cleared on all nodes following CAS-1 handshake.

This multicast ALE link **may** also be used for 1:1 traffic between the participating nodes. In general, it will be preferable to terminate the multicast link when there is no further traffic and to use 1:1 ALE links for 1:1 traffic.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

B.5.4.   Duplex with 1:1 ALE

In order to use Duplex transfer between a pair of nodes without ALE, it is necessary to configure a fixed frequency for each direction of transfer between a pair of nodes.

When a pair of nodes each have duplex support and ALE is used, duplex communication can be negotiated.   In order to do this, a node shall determine if the peer can support duplex communication. There are two methods to achieve this:

1. A priori knowledge; or

2. The Data Transfer Sublayer enables determination that a peer node is able to support duplex by use of EOWs.

If a node determines that both peers support duplex, duplex communication **may** be negotiated using ALE.   It is recommended that this is done. On receiving an ALE link, a node with a duplex peer **may** use ALE to negotiate a second connection with the peer.   If this is successfully done, the first ALE link **shall** be used by the node which initiated it to send data and the second ALE link **shall** be used to send data in the reverse direction.

B.6.   Channel Access Protocol Type 1 and C_PDUs

Where ALE is not used, the local node will be connected to a channel, which can be fixed frequency or with variable frequency determined by an external process (e.g., selection of frequency based on a fixed schedule). Other nodes may be connected to the same channel.

The local node **may** know the addresses of all other nodes that can be connected to the channel. If this is the case and data is addressed to another node, it **shall** be rejected with reason code "Address Not Known".

The co-ordination of the making and breaking of Physical Links (hereinafter referred to as 'CAS 1 Linking Protocol') between two nodes **shall** be performed solely by the Channel Access Sublayer. If ALE is not used, the CAS 1 linking protocol **shall** be used.

If the CAS 1 Linking Protocol is omitted (when ALE is used) and the response to ARQ-data DPDUs is a Warning DPDU indicating that a connection is not made, then the CAS 1 Linking Protocol **shall** be followed and the data resent.

If a slave node (receiving station) accepts an ALE link and a CAS-1 Physical-Link Request C_PDU is received, then the receiving node **shall** respond in accordance with the requirements of the 'CAS 1 Linking Protocol', accepting or rejecting the request as is appropriate for its state.

When ALE is used, the CAS 1 linking protocol **shall** respond to the reception of a type 4 PHYSICAL_LINK_BREAK C_PDU by sending a type 5 PHYSICAL_LINK_BREAK_CONFIRM C_PDU as specified in Section B.4.2.2. The ALE link is then closed and the physical link **shall** be broken at that point.

A Node shall use a Physical Link to support control and data exchange for Soft-Link Data Exchange Sessions

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

as requested by the Subnetwork Interface Sublayer.

A Node **may** have Physical Links with more than one other node at a time, one Physical Link per remote node; i.e., a Node may "Make" a new Physical Link with another node before it "Breaks" any Physical links.

There **shall** be no explicit peer-to-peer communication required to switch from use of one Physical Link to another. When 1:1 ALE is used, the Channel Access sublayer enforces that there is at most one active ALE link per channel and that for each active ALE link there is a single physical link associated with that ALE link.

The Channel Access Sublayer CAS 1 linking protocol **shall** communicate with peer sublayers in other nodes using the protocols defined here in order to:

1. Make and break physical links

2. Deliver S_PDUs between Subnetwork Interface Sublayers at the local node and remote node(s).

B.6.1. Type 1 Channel Access Sublayer Data Protocol Units (C_PDUs)

The following C_PDUs **shall** be used for peer-to-peer communication between Channel Access Sublayers in the local and remote node(s):

C_PDU NAME	*Type Code*
DATA C_PDU	TYPE 0
PHYSICAL LINK REQUEST	TYPE 1
PHYSICAL LINK ACCEPTED	TYPE 2
PHYSICAL LINK REJECTED	TYPE 3
PHYSICAL LINK BREAK	TYPE 4
PHYSICAL LINK BREAK CONFIRM	TYPE 5

The first argument and encoded field of all C_PDUs **shall** be the C_PDU Type

The remaining format and content of these C_PDUs **shall** be as specified in the subsections that follow.

Unless noted otherwise, argument values encoded in the C_PDU bit-fields **shall** be mapped into the fields in accordance with CCITT V.42, 8.1.2.3, i.e.:
  1. when a field is contained within a single octet, the lowest bit number of the field **shall** represent the lowest-order (i.e., least-significant-bit) value;
  2. when a field spans more than one octet, the order of bit values within each octet **shall** decrease progressively as the octet number increases. The lowest bit number associated with the field **shall** represent the lowest-order (i.e., least significant bit) value.

Unless noted otherwise, bit-fields specified as NOT USED **shall** be encoded with the value '0' (i.e., zero).

B.6.1.1. DATA C_PDU

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

Type :

"0" = DATA C_PDU

Encoding :

C_PDU TYPE 0 (DATA)

C_PCI	S_PDU

TYPE
(= 0)

(a)   Generic Encoding

7	6	5	4	3	2	1	0
TYPE = 0				NOT USED			
0	0	0	0				
S_PDU							

(b)   Bit-Field Map

Figure B-1. Generic Encoding and Bit-Field Map of the DATA C_PDU

Description :

The DATA (TYPE 0) C_PDU **shall** be used to send an encapsulated S_PDU from the local node to a remote node.

The *Type* argument **shall** be encoded in the first four-bit field of the DATA C_PDU as shown in Figure B-1. The value of the *Type* argument for the DATA C_PDU **shall** be zero.

The remaining octets of the DATA C_PDU **shall** contain the encapsulated S_PDU and only the encapsulated S_PDU.

For the Channel Access sublayer request to the lower layers of the subnetwork to deliver a C_PDU, the delivery service requirements for a DATA C_PDU **shall** be the same as the S_PDU that it contains, i.e.:

a)   C_PDUs **shall** be sent using the normal Data Delivery service provided by the lower sublayer;

b)   the DELIVERY mode specified by the Subnetwork Interface Sublayer for the encapsulated S_PDU (i.e., ARQ, non-ARQ, etc.) also **shall** be assigned to the C_PDU by the Channel Access sublayer as the delivery mode to be provided by the lower sublayer.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

B.6.1.2. PHYSICAL LINK REQUEST C_PDU

Type :
     "1" = PHYSICAL LINK REQUEST

Encoding :

MSB							LSB
7	6	5	4	3	2	1	0
Type=1				Not Used	Short	Ed4	Not Used

Figure B-2. Generic Encoding and Bit-Field Map of the PHYSICAL LINK REQUEST C_PDU

Description :
     The PHYSICAL LINK REQUEST C_PDU **shall** be transmitted by a Channel Access sublayer to request the making of the Physical Link.

     The PHYSICAL LINK REQUEST C_PDU **shall** consist of the arguments *Type* and *Link*.

     The value of the *Type* argument for the PHYSICAL LINK REQUEST C_PDU **shall** be '1' (i.e., one), encoded as a four-bit field as shown in Figure B-2.

     The bits not used in the encoding of the PHYSICAL LINK REQUEST C_PDU **shall** be reserved for future use and not used by any implementation.

     The Ed4 bit indicates to the peer that the local node support edition 4 (or subsequent version) of the protocol. If the peer is known to support edition 4 (or subsequent) this bit **shall** be set. If the peer is known to not support edition 4 (or subsequent) this bit **shall** not be set. Otherwise this bit **may** be set in order to discover if the peer supports edition 4.

     If the Ed4 bit is set, the Short bit **may** be set. This indicates a request to use D_PDUs with short LFSN, which can give performance advantage at very low speed.

     When a PHYSICAL LINK REQUEST C_PDU is transmitted the local Node is not linked to another Node and therefore the ARQ transmission mode is not supportable by the Data Transfer Sublayer. Therefore, the PHYSICAL LINK REQUEST C_PDU **shall** be sent by the Channel Access Sublayer requesting the lower-layer Expedited Non-ARQ Data Service (i.e., transmission using Type 8 D_PDUs) in accordance with the Annex C specification of the Data Transfer Sublayer.

     A Channel Access sublayer which receives a PHYSICAL LINK REQUEST C_PDU **shall**

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

respond with either a PHYSICAL LINK ACCEPTED (TYPE 2) C_PDU or a PHYSICAL LINK REJECTED (TYPE 3) C_PDU, as appropriate.

B.6.1.3. PHYSICAL LINK ACCEPTED C_PDU

Type :
      "2" = PHYSICAL LINK ACCEPTED

Encoding :

	MSB 7	6	5	4	3	2	1	LSB 0
0		Type=1			Not Used	Short	Ed4	Not Used

Figure B-3. Generic Encoding and Bit-Field Map of the PHYSICAL LINK ACCEPTED C_PDU

Description :

The PHYSICAL LINK ACCEPTED (TYPE 2) C_PDU **shall** be transmitted by a peer sublayer as a positive response to the reception of a TYPE 1 C_PDU (PHYSICAL LINK REQUEST).

PHYSICAL LINK ACCEPTED (TYPE 2) C_PDU **shall** consist only of the argument *Type*. [Note: A Node can never request the establishment of more than one Physical Link with another given node, therefore, the request for which this message is a response is uniquely identified by the node address of the node to which the PHYSICAL LINK ACCEPTED C_PDU is sent.]

The *Type* argument **shall** be encoded as a four-bit field containing the binary value 'two' as shown in Figure B-3.

The bits not used in the encoding of the PHYSICAL LINK REQUEST C_PDU **shall** be reserved for future use and not used by any implementation.

The Ed4 bit indicates to the peer that the local node support edition 4 (or subsequent version) of the protocol. If the equivalent bit was set in the PHYSICAL LINK REQUEST C_PDU which is being responded to this bit **shall** be set. Otherwise, this bit **shall not** be set.

If the Ed4 bit is set and the equivalent bit was set in the PHYSICAL LINK REQUEST C_PDU which is being responded, the Short bit **may** be set. This indicates that D_PDUs with short LFSN, which can give performance advantage at very low speed **shall** be used with this link.

When a PHYSICAL LINK ACCEPTED (TYPE 2) C_PDU is transmitted the local Node is not linked to another Node and therefore the ARQ transmission mode is not supportable by the Data Transfer Sublayer. Therefore, the PHYSICAL LINK ACCEPTED C_PDU **shall** be sent by the Channel Access Sublayer requesting the lower-layer's Expedited Non-ARQ Data Service (i.e., transmission using Type 8 D_PDUs) in accordance with the Annex C specification of the Data

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

Transfer Sublayer.

B.6.1.4. PHYSICAL LINK REJECTED C_PDU

Type :
"3" = PHYSICAL LINK REJECTED

Encoding :

TYPE (=3)	REASON

7	6	5	4	3	2	1	0
TYPE = 3₀ 0				REASON MSB -- -- LSB			

Figure B-4. Generic Encoding and Bit-Field Map of the PHYSICAL LINK REJECTED C_PDU

Description :

The PHYSICAL LINK REJECTED (TYPE 3) C_PDU **shall** be transmitted by a peer sublayer as a negative response to the reception of a TYPE 1 C_PDU (PHYSICAL LINK REQUEST).

The PHYSICAL LINK REJECTED (TYPE 3) C_PDU **shall** consist of two arguments: *Type* and *Reason.*

The *Type* argument **shall** be encoded as a four-bit field containing the binary value 'three' as shown in Figure B-4.

The *Reason* argument **shall** be encoded in accordance with Figure B-4 and the following table:

Reason	Value
Reason Unknown	0
Broadcast-Only-Node	1
Not Used	3
Not Used	4
Too many active links	5
unspecified	6-15

The value 0, indicating an unknown reason for rejecting the Make request for the physical link, is always a valid reason for rejection. Reasons corresponding to values in the range 4-15 are currently unspecified and unused in the STANAG.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

When a PHYSICAL LINK REJECTED C_PDU is transmitted the local Node is not linked to another Node and therefore the ARQ transmission mode is not supportable by the Data Transfer

Note that "Link busy" is not included in the list of possible reasons for rejecting anode's request to establish a link since this function is provided by the WARNING frame type of the Data Transfer Sublayer, as specified in Annex C.

Sublayer. Therefore, the PHYSICAL LINK REJECTED C_PDU **shall** be sent by the Channel Access Sublayer requesting the lower-layer Expedited Non-ARQ Data Service (i.e., transmission using Type 8 D_PDUs) in accordance with the Annex C specification of the Data Transfer Sublayer.

B.6.1.5. PHYSICAL LINK BREAK C_PDU

Type :
    "4" = PHYSICAL LINK BREAK
Encoding :

TYPE (=4)	REASON

| 0	7	6	5	4	3	2	1	0
		TYPE = 4				REASON		
	0	1	0	0	MSB	--	--	LSB

Figure B-5. Generic Encoding and Bit-Field Map of the PHYSICAL LINK BREAK C_PDU

Description :
    The PHYSICAL LINK BREAK C_PDU **shall** be transmitted by either of the peer Channel Access sublayers involved in an active Physical Link to request the breaking of the link.

The PHYSICAL LINK BREAK C_PDU **shall** consists of two arguments: *Type and Reason.*

The *Type* argument **shall** be encoded as a four-bit field containing the binary value 'four' as shown in Figure B-5.

The *Reason* argument **shall** be encoded in accordance with Figure B-5 and the following table:

Reason	Value

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

Reason Unknown	0
Higher-Layer-Request	1
Switching to Broadcast- Data-	2
Higher-Priority-Link-Request-Pending	3
No More Data	4
unspecified	4-15

The value 0, indicating an unknown reason for requesting the breaking of a physical link, is always a valid reason. Reasons corresponding to values in the range 4-15 currently are not defined in this STANAG.

A peer sublayer which receives the PHYSICAL LINK BREAK C_PDU **shall** immediately declare the Physical Link as broken and respond with a PHYSICAL LINK BREAK (TYPE 5) C_PDU as specified in section B.3.1.6 below.

The PHYSICAL LINK BREAK C_PDU **shall** be sent by the Channel Access Sublayer requesting the lower-layer Expedited Non-ARQ Data Service (i.e., transmission using Type 8 D_PDUs) in accordance with the Annex C specification of the Data Transfer Sublayer.
[Note: The reason the PHYSICAL LINK BREAK C_PDU is sent with the non-ARQ data service, even though a physical link exists at the time that it is sent, is because the receiving peer will immediately declare the Link as broken. The receiving peer will therefore not have time to send ARQ acknowledgements for the C_PDU.]

B.6.1.6. PHYSICAL LINK BREAK CONFIRM C_PDU

Type :
    "5" = PHYSICAL LINK BREAK CONFIRM
Encoding :

Figure B-6. Generic Encoding and Bit-Field Map of the PHYSICAL LINK BREAK CONFIRM C_PDU

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

Description :

The PHYSICAL LINK BREAK CONFIRM (TYPE 5) C_PDU **shall** be transmitted by a Channel Access sublayer as a response to a TYPE 4 "PHYSICAL LINK BREAK" C_PDU.

The PHYSICAL LINK BREAK CONFIRM (TYPE 5) C_PDU **shall** consist of the single argument *Type*.

The *Type* argument **shall** be encoded as a four-bit field containing the binary value 'five' as shown in Figure B-6.

The PHYSICAL LINK BREAK CONFIRM (TYPE 5) C_PDU **shall** be sent by the Channel Access Sublayer requesting the lower-layer Expedited Non-ARQ Data Service (i.e., transmission using Type 8 D_PDUs) in accordance with the Annex C specification of the Data Transfer Sublayer.

Upon receiving a PHYSICAL LINK BREAK CONFIRM (TYPE 5) C_PDU, the peer which initiated the breaking of the Link **shall** declare the Link as broken.

B.6.2. Type 1 Channel Access Sublayer Peer-to-Peer Communication Protocol

Requirements on the Type 1 Channel Access Sublayer Peer-to-Peer Communication Protocols are specified in this section and the subsections below.

The Channel Access sublayer **shall** perform all peer-to-peer communications for protocol control using the following C_PDUs:

C_PDUs for Peer-to-Peer Protocol Control	Type Code
PHYSICAL LINK REQUEST	TYPE 1
PHYSICAL LINK ACCEPTED	TYPE 2
PHYSICAL LINK REJECTED	TYPE 3
PHYSICAL LINK BREAK	TYPE 4
PHYSICAL LINK BREAK CONFIRM	TYPE 5

All peer-to-peer communications **shall** be done by the Channel Access sublayer requesting the lower-layer Expedited Non-ARQ Data Service (i.e., transmission using Type 8 D_PDUs in accordance with the Annex C specification of the Data Transfer Sublayer).

The criteria for accepting or rejecting Physical Links Requests **shall** be as follows. A request to establish a Physical Link **shall** be accepted in Single Peer Access mode, as receipt of the request implies that the 1:1 underlying channel is open. In Multiple Simultaneous Peer Access mode, the Physical Link **shall** be accepted, unless the number of active physical links exceeds a configurable limit or another operational reason.

B.6.2.1. Protocol for Making a Physical Link

In the specification that follows of the protocol for making a physical link, the node which requests the physical

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

link is referred to as the Caller or Calling node, while the node which receives the request will be referred to as the Called node.

The protocol for making the physical link **shall** consist of the following steps:

Step 1-Caller:
a) The Calling Node's Channel Access Sublayer **shall** send a PHYSICAL LINK REQUEST (TYPE 1) C_PDU to initiate the protocol,
b) Upon sending the PHYSICAL LINK REQUEST (TYPE 1) C_PDU the Channel Access Sublayer **shall** start a timer which is set to a value greater than or equal to the maximum time required by the Called Node to send its response (this time depends on the modem parameters, number of re-transmissions, etc.),
c) The maximum time to wait for a response to a PHYSICAL LINK REQUEST (TYPE 1) C_PDU shall be a configurable parameter in the implementation of the protocol.

Step 2-Called:
a) On receiving a PHYSICAL LINK REQUEST (TYPE 1) C_PDU, a Called node **shall** determine whether or not it can accept or reject the request as follows:
    (1) If the Called node has the maximum number of active Physical Links, the Called node **shall** reject the request with reason "Too many active links",
    (2) otherwise, the Called node **shall** accept the request for a Physical Link.
b) After determining if it can accept or reject the PHYSICAL LINK REQUEST, a Called node **shall** respond as follows:
    (1) if a Called node accepts the physical link request, it **shall** respond with a PHYSICAL LINK ACCEPTED (TYPE 2) C_PDU,
    (2) otherwise, when a Called node rejects the physical link request, it **shall** respond with a PHYSICAL LINK REJECTED (TYPE 3) C_PDU.
b) After a PHYSICAL LINK ACCEPTED (TYPE 2) C_PDU is sent, the called channel access sublayer **shall** declare the physical link made and transition to a state in which it executes the protocol for data exchange using C_PDUs.
c) If further PHYSICAL LINK REQUEST (TYPE 1) C_PDUs are received from the same address after the link is made, the Channel Access sublayer **shall** again reply with a PHYSICAL LINK ACCEPTED (TYPE 2) C_PDU.
d) If at least one DATA (TYPE 0) C_PDU is not received on a newly activated Physical Link after waiting for a specified maximum period of time, the Called Node **shall** abort the Physical Link and declare it inactive.
e) The maximum period of time to wait for the first DATA (TYPE 0) C_PDU before aborting a newly activated Physical Link **shall** be a configurable parameter in the implementation.

Step 3. Caller

There are two possible outcomes to the protocol for making a physical link: success or failure:
a) Upon receiving a PHYSICAL LINK ACCEPTED (TYPE 2) C_PDU, the calling Channel Access Sublayer **shall** proceed as follows:
    (1) the Calling node **shall** declare the Physical Link as *successfully Made*, otherwise,
b) upon receiving a PHYSICAL LINK REJECTED (TYPE 3) C_PDU, the Channel Access Sublayer **shall** declare the Physical Link as *Failed,* otherwise,
c) upon expiration of its timer without any response having been received from the remote node, the Channel Access Sublayer **shall** repeat Step 1 (i.e. send a PHYSICAL LINK REQUEST (TYPE 1) C_PDU and set a response time) and await again a response from the remote node.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

The maximum number of times the Caller sends the PHYSICAL LINK REQUEST (TYPE 1) C_PDU without a response from the called node of either kind **shall** be a configurable parameter in the implementation of the protocol.

After having repeated Step 1 the configurable maximum number of times without any response having been received from the remote node, the Caller's Channel Access Sublayer **shall** declare the protocol to make the physical link as *Failed*.

| Figure B-7 (a) and (b) show the nominal procedures followed by the Caller and Called Channel Access Sublayers for Making a Physical Link. This STANAG acknowledges that other implementations may meet the stated requirements.
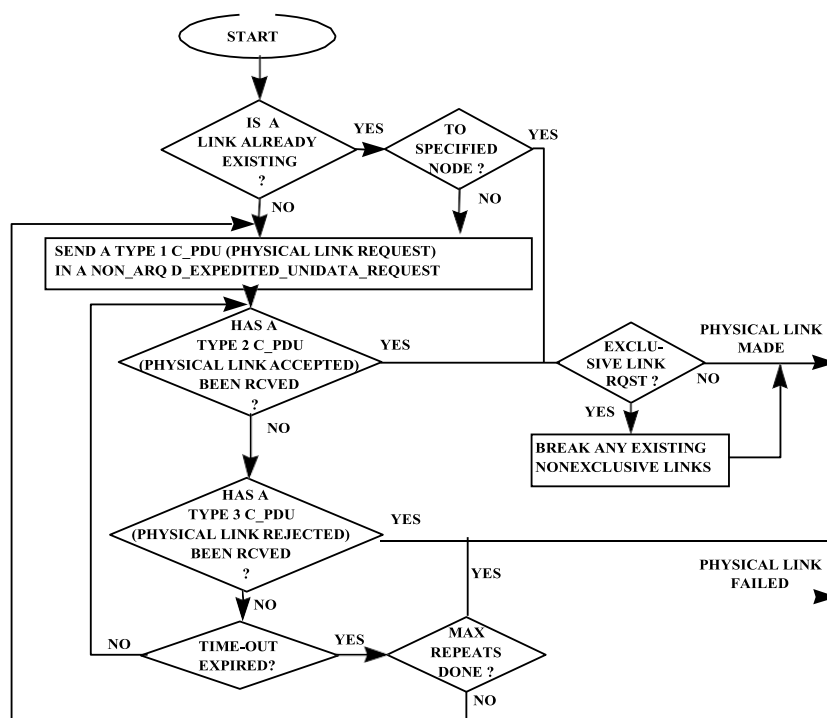
Figure B-7 (a). Type 1 Channel Access Sublayer protocol for Making a Physical Link (Caller Peer)

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4
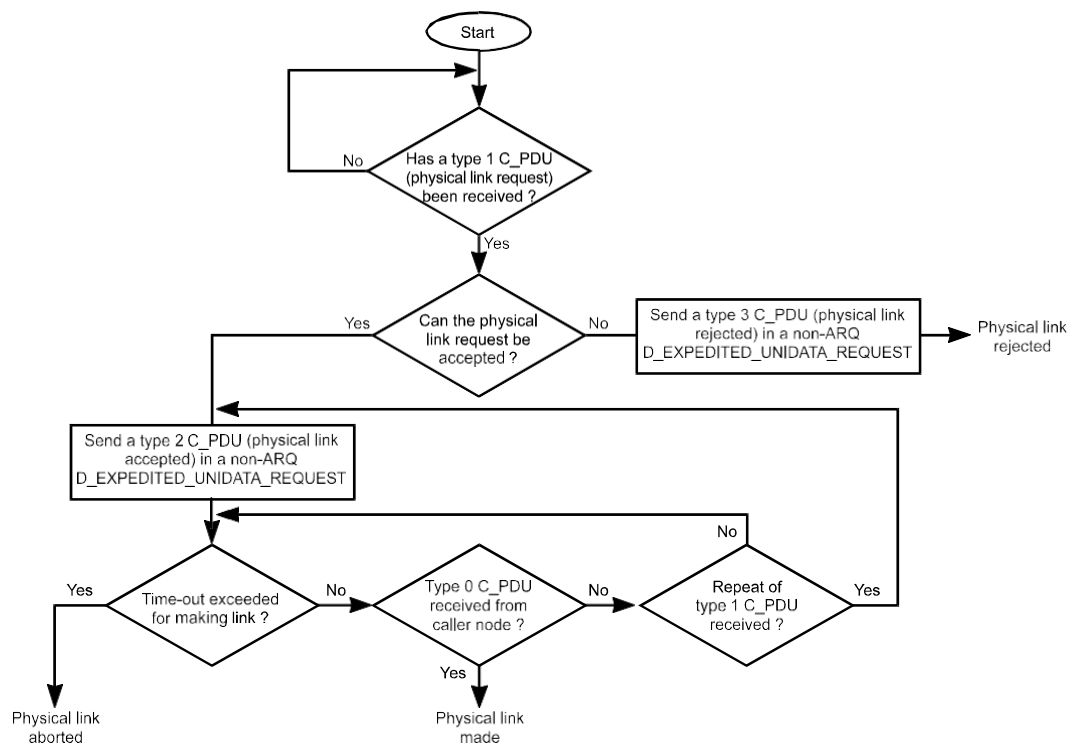
Figure B-7. Type 1 Channel Access Sublayer Protocol for Making a Physical Link (Called Peer)

After having declared the Physical Link as successfully made, the peer Channel Access Sublayers do not carry out any further handshake in order to confirm the successful operation of the Link.

B.6.2.2. Determining Edition 4 Support

The physical link setup is also used to determine support for this edition of the protocol (or subsequent editions). On receipt of a PHYSICAL LINK REQUEST C_PDU with the Ed4 bit set, a node can determine that a peer supports edition 4.

It is likely that an edition 3 (or earlier) implementation receiving a PHYSICAL LINK REQUEST C_PDU with the Ed4 bit set will ignore this bit and respond with a PHYSICAL LINK ACCEPT C_PDU which does not have the Ed4 bit set.   This will enable the local node to determine that the peer does not support edition 4. This mechanism provides a robust way to determine if a peer supports edition 4 (or subsequent).

It is possible, but unlikely, that the edition 3 peer will respond to a C_PDU with the Ed4 bit set by treating it as an error and responding with a WARNING D_PDU.   In the event that this happens the local peer **shall** record that the peer does not support edition 4 and **shall** repeat the link establishment without the Ed4 bit set. For subsequent links, the record of edition support can avoid repeating the failed link attempt.

When both peers support edition 4 (or subsequent) the default DTS behavior will be to use the family of D_PDUs with 16 bit Frame Sequence Number.   Use of the Short bit allows peers to negotiate use of D_PDUs with 8 bit Frame Sequence Number. This could be beneficial in scenarios where it is known that speeds will bet 1200 bps or less, as the D_PDU header overhead is slightly reduced.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

B.6.2.3. Protocol for Breaking a Physical Link

In the specification that follows of the protocol for breaking a physical link, the node which requests the breaking of the Physical Link will be referred to as the Initiator or initiating node, while the node which receives the request will be referred to as the Responder or responding node. Physical links shall always be explicitly broken, even if set up implicitly by use of ALE.

Physical Links will be broken in the following situations:

1. No more traffic to be sent over the link.

2. In Single Peer Access Model:

    a. Higher priority traffic has arrived for a different peer, and the link needs to be closed as soon as possible.

    b. A link is closed to provide fair access between peers.

3. Operator chooses to break a link, perhaps due to excess queued traffic.

It is recommended that SIS client flow control is used to avoid large queue build up in the DTS. However, there needs to be some measure of queueing in DTS in order for the DTS to operate efficiently. When a link is broken prior to all traffic being sent, it is generally desirable to close the link quickly. For this reason, link break procedure makes use of expedited data, which will ensure that the link break protocol is not blocked behind other traffic on the link.

The protocol for breaking the Physical Link **shall** consist of the following steps:

Step 1: Initiator

a)   To start the protocol, the Initiator's Channel Access Sublayer **shall** send a type 4 C_PDU (PHYSICAL LINK BREAK).

b)   Upon sending this C_PDU the Channel Access Sublayer **shall** start a timer which is set to a value greater than or equal to the maximum time required by the Called Node to send its response (this time depends on the modem parameters, number of re-transmissions, etc.).

Step 2: Responder

c)   Upon receiving the type 4 C_PDU the Responder's Channel Access Sublayer **shall** declare the Physical link as *broken* and send a PHYSICAL LINK BREAK CONFIRM (TYPE 5) C_PDU.

Step 3: Initiator

d)   Upon receiving a PHYSICAL LINK BREAK CONFIRM (TYPE 5) C_PDU, the Initiator's Channel Access Sublayer **shall** declare the Physical Link as *broken*.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

e)   Upon expiration of its timer without any response having been received from the remote node, the Initiator's Channel Access Sublayer **shall** repeat step 1 and wait again for a response from the remote node.

f)   After having repeated Step 1 a maximum number of times (left as a configuration parameter) without any response having been received from the remote node, the Initiator's Channel Access Sublayer **shall** declare the Physical Link as *broken*.

Figures B-8 (a) and (b) show the procedures followed by the Initiator and Responding Channel Access Sublayers for Breaking a Physical Link.
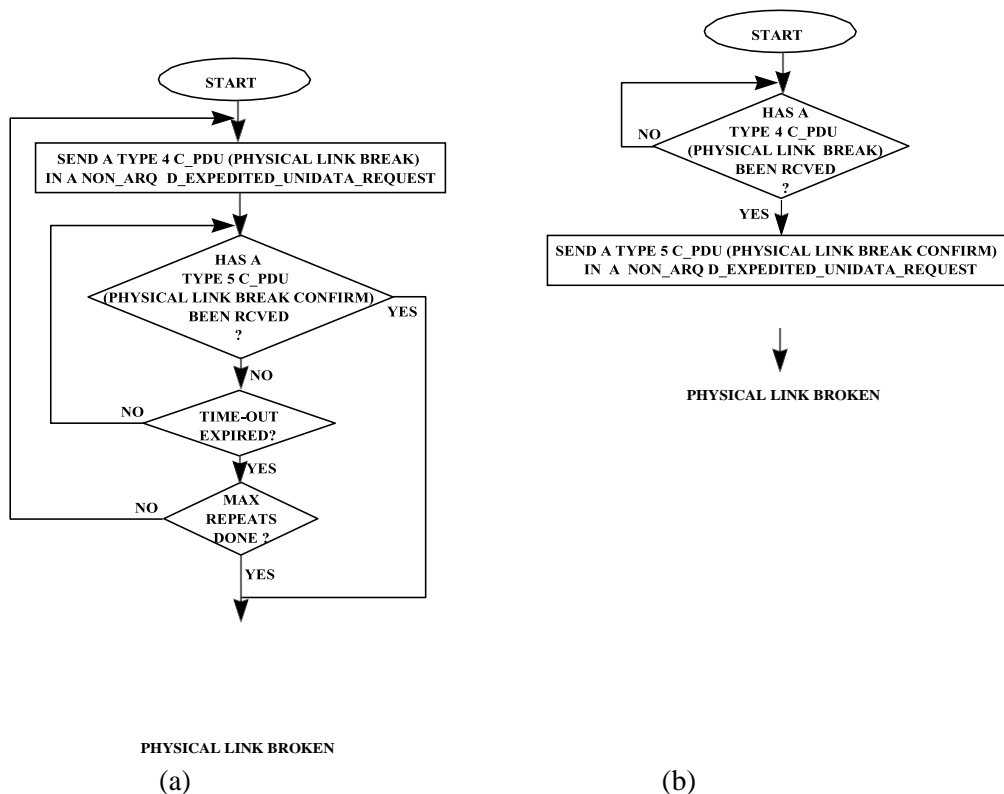
**Figure B-8. Type 1 Channel Access Sublayer protocol for Breaking a Physical Link
(a) Initiator Peer. (b) Responding Peer.**

ANNEX B
STANAG 5066
Edition 3

B.7.   <u>Protocol for Exchanging Data C_PDUs.</u>

The protocol for exchanging C_PDUs between peers is simple and straightforward, since the Channel Access Sublayer relies entirely on its supporting (i.e, Data Transfer) sublayer to provide delivery services.

The sending peer **shall** accept S_PDUs from the Subnetwork Interface Sublayer, envelop them in a "DATA" C_PDU (by adding the C_PCI) and send them to its receiving peer via its interface to the Data Transfer

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX B
STANAG 5066
Edition 4

Sublayer.

The receiving peer **shall** receive DATA C_PDUs from the Data Transfer Sublayer interface, check them for validity, strip off the C_PCI, and deliver the enveloped S_PDU to the Subnetwork Interface Sublayer.

B.8.   Changes in Edition 4

This annex has the following changes relative to Edition 3.

1. Explicit setting out of single peers access vs multiple peer access as part of the channel model. This was implicit in Edition 3. It needs to be made explicit in order to clearly document ALE and Duplex operation.

2. Use of 1:1 ALE for connecting to a peer is documented explicitly.

3. The model of duplex operation is set out, both with and without ALE.

4. Where ALE is not used, the specification requires CAS-1 linking. The generic text around handling ALE-like systems is removed.

5. The CAS-1 negotiation is extended to determine if the peer supports Edition 4, and if so to negotiate the choice of D_PDUs to use on the soft link.

6. Some new Error types are introduced, arising from ALE and Duplex support.

7. C_CHANNEL_AVAILABILITY service primitive is removed. It is listed, but not defined. It does not appear to be useful.

8. Errors for C_ primitives explicitly listed.

9. Support for hard links removed.

10. Incorporation of notes from STANAG 5066 Edition 3 Annex H.

11. Removal of handling Expedited Data, in line with Annex A.

**Annex C: Data Transfer Sublayer
(Mandatory)
DRAFT 1.3 of 9th April 2020**

The Data Transfer Sublayer is responsible for the efficient transfer across the radio link of protocol data units from the Management Sublayer (i.e., M_PDUs) and Channel Access Sublayer (i.e., C_PDUs) in the STANAG 5066 profile. Usually, but not always, this means the error free delivery of C_PDUs to the Channel Access Sublayer. Efficient transmission over the radio channel of protocol data units of the Data Transfer Sublayer (i.e., D_PDUs, which contain the C_PDUs or M_PDUs) is achieved principally by two mechanisms. The first is segmentation of the large C_PDUs into smaller D_PDUs if necessary and the second is the selective repetition (selective ARQ) by the sending node of D_PDUs which were received in error. The other mode (non ARQ) of transmitting involves the segmentation of the large C_PDUs into smaller D_PDUs and combining the received D_PDUs such that the segmented C_PDU can be reconstructed in a "best-effort" attempt.

C.1.   Data Transfer Sublayer Service Definition

Depending on the application and service-type requested by higher sublayers, the user service provided by the Data Transfer Sublayer **shall** [(1)] be either a simple **non ARQ service** or a reliable **selective ARQ service**, as specified herein.

The Data Transfer Sublayer **shall** [(2)] provide "sub-modes" for **non ARQ** and reliable **selective ARQ** delivery services, which influence the characteristics of the particular service, as specified below

In addition to the Selective ARQ and Non-ARQ services provided to the upper sublayers, the Data Transfer Sublayer shall provide an Idle Repeat Request service for peer-to-peer communication with the Data Transfer Sublayer of other nodes.

C.1.1.   Non-ARQ Service

In the **non ARQ service** error-check bits (i.e., cyclic-redundancy-check or CRC bits) applied to the D_DPU **shall** [(1)] be used to detect errors, and any D_PDUs that are found to contain transmission errors **shall** [(2)] be discarded by the data transfer sublayer protocol entity, except as noted below in support of delivering partial C_PDUs.

In the **non ARQ** mode, the following submodes may be specified:
- regular data service.
- expedited data service.

"in order" delivery of C_PDUs is not guaranteed and C_PDUs **shall** be delivered in the order in which they arrive. "in-order" delivery is a service for ARQ only. There is some implication in Edition 3 that this service is available for non-ARQ. Experience has shown that attempting to provide this service leads to operational problems.

ANNEX C
STANAG 5066
Edition 4

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

Delivery of complete and error free C_PDUs is not guaranteed. Delivery of C_PDUs is not guaranteed. A special mode of the non-ARQ service **shall** [3] be available to reconstruct partial C_PDUs from D_PDUs in error and deliver the partial and potentially erroneous elements to the Channel Access Sublayer.

C.1.2.   Selective ARQ Service

The reliable Selective ARQ service shall (1) use CRC check bits and flow control procedures, such as requests for retransmission of D_PDUs in which errors have been detected, to provide a reliable data transfer service.

Transfer of complete and error free C_PDUs is guaranteed.

In the **Selective ARQ** service, the following submodes may be specified:

- regular data service.
- expedited data service.

Both regular and expedited services offer delivery confirmation for data delivery to the peer node.

Both regular and expedited services offer an "in-order" option, which will deliver C_PDUs to the remote peer and associated channel access service in the order they were submitted. If this service is not selected, C_PDUs will be delivered as they arrive.

NOTE: "in order" is implemented at DTS level and so will apply to all SAPs using it. This can lead to one SAP blocking another. It is **recommended** to avoid use of "in order" whenever possible.

C.1.3.   Queueing and Flow Control

It is expected that the DTS will queue C_PDUs in order to effectively manage transmission, but that the amount of data queued will be limited.   The DTS will provide flow control information to the CAS, which will be used by the SIS to limit data being provided to the DTS.

Experience with DTS suggests that DTS queue lengths should be kept as short as possible, consistent with efficient operation of the DTS. Use of SIS flow control allows better application level monitoring and better behavior in error situations.

C.1.4.   Full Duplex Operation

Annex C can operate with two types of configuration:

1. Single Channel. Here D_PDUs are all sent to a channel which is shared by multiple nodes with access to the channel controlled by the MAC layer specified in Annex J.   D_PDUs may be sent to multiple nodes, optionally using half duplex communication with each of the peer nodes.

2. Full Duplex, using two channels (Receive and Transmit) to communicate with a single peer. In Full

Duplex mode, the CAS **shall** ensure that the DTS is used only for communications with a single peer.

The core operations and state machine operation of the DTS are common to both single channel and full duplex. The details of transmission are different, and the rules for both types of operation are set out in Section C.7.5.

C.2.   Interface Primitives Exchanged with the Channel Access Sublayer

The implementation of the interface between the Data Transfer Sublayer and the Channel Access Sublayer is not mandated or specified by this STANAG. Since the interface is internal to the subnetwork architecture and may be implemented in a number of ways it is considered beyond the scope of STANAG 5066. A model of the interface has been assumed, however, for the purposes of discussion and specification of other sublayer functions.

Despite the advisory nature of the conceptual model of the internal interface between the Data Transfer Sublayer and the Channel Access Sublayer, there are some mandatory requirements that are placed on any interface implementation.

The interface must support the service-definition for the Data Transfer Sublayer, i.e.:

1. The interface **shall** [1] allow the Channel Access Sublayer to submit protocol data units (i.e., C_PDUs) for transmission using the regular and expedited delivery services provided by the Data Transfer Sublayer.

2. The interface **shall** [2] allow the Data Transfer Sublayer to deliver C_PDUs to the Channel Access Sublayer.

3. The interface **shall** [3] permit the Channel Access Sublayer to specify the delivery services, priority, and time-to-die required by the C_PDUs when it submits them to the Data Transfer Sublayer.

4. The interface **shall** [4] permit the Data Transfer Sublayer to specify the delivery services that were used by received C_PDUs when it submits them to the Channel Access Sublayer.

5. The interface **shall** [5] permit the Channel Access Sublayer to specify the destination address to which C_PDUs are to be sent.

6. The interface **shall** [6] permit the Data Transfer Sublayer to specify the source address from which C_PDUs are received, and the destination address to which they had been sent.

7. The interface **shall** [7] permit the Data Transfer Sublayer to notify the Channel Access sublayer when a warning indication (i.e., a WARNING D_PDU) has been received from a remote peer, the source and destination address associated with the warning, the reason for the warning, and the event (i.e., message type) that triggered the warning message.

8. The interface **shall** [8] permit the Data Transfer Sublayer to notify the Channel Access sublayer that a warning indication (i.e., a WARNING D_PDU) has been sent to a remote peer, the destination address associated with the warning, the reason for the warning, and the event (i.e., message type) that triggered the warning message.

9. The interface **shall** [9] permit the Channel Access sublayer to notify the Data Transfer Sublayer that a Link has been established with a given node.

10. The interface **shall** [10] permit the Channel Access sublayer to notify the Data Transfer Sublayer that a Link has been terminated with a given node.

11. The interface **shall** [11] permit the Data Transfer Sublayer to notify the Channel Access sublayer that a Link has been lost with a given node.

Additionally, the protocol-control information from the Channel Access sublayer that is required for the management of the Data Transfer Sublayer **shall** [12] not be derived from knowledge of the contents or format of any client data or U_PDUs encapsulated within the C_PDUs exchanged over the interface.
[Note: user's that encrypt their traffic prior to submittal may use the subnetwork. Subnetwork operation must be possible with client data in arbitrary formats that are unknown to the subnetwork, therefore any service requirements or indications must be provided by interface control information provided explicitly with the user data.]

The interface **may** use knowledge of the contents of C_PDUs (excluding the contents of any encapsulated U_PDUs) to derive protocol control information for the Data Transfer sublayer. This approach is highly discouraged, however. The recommended approach for implementation is that information required for protocol control within the Data Transfer sublayer should be provided explicitly in appropriate interface primitives.

In keeping with accepted practice in the definition of layered protocols, and as a means for specifying the operations of the sublayers that are mandated by this STANAG, the communication between the Data Transfer Sublayer and the Channel Access Sublayer is described herein with respect to a set of Primitives. The interface Primitives are a set of messages for communication and control of the interface and service requests made between the two layers.

By analogy to the design of the client-subnetwork interface specified in Annex A, the technical specification of the Data Transfer Sublayer assumes communication with the Channel Access Sublayer using primitives prefixed with a "D_". A minimal set of D_Primitives has been assumed that meet the requirements stated above and the general function for each D_Primitive is given in Table C-1. These D_Primitives are given without benefit of a list of arguments or detailed description of their use.

Table C-1 – Nominal Definition of D_Primitives for the Interface between the Data Transfer Sublayer and the Channel Access sublayer
(non-mandatory, for information-only)

NAME OF PRIMITIVE	DIRECTION (see Note)	COMMENTS
D_UNIDATA_REQUEST	CAS →DTS	Request to send an encapsulated C_PDU using the regular delivery service to a specified destination node address with given priority, time-to-die, and delivery mode.
D_UNIDATA_REQUEST_CONFIRM	DTS →CAS	Confirmation that a given C_PDU has been sent using the regular delivery service
D_UNIDATA_REQUEST_REJECTED	DTS →CAS	Notification that a given C_PDU could not be sent using the regular delivery service and the reason why it was rejected by Data transfer Sublayer.

D_UNIDATA_INDICATION	DTS →CAS	Delivers a C_PDU that has been received using the regular delivery service from a remote node, with additional indications of the transmission service given to the C_PDU, the source node address, and the destination node address (e.g., if a group address is the destination).
D_EXPEDITED_UNIDATA_REQUEST	CAS →DTS	Request to send an encapsulated C_PDU using the expedited delivery service to a specified destination node address with a specified delivery mode.
D_EXPEDITED_UNIDATA_ REQUEST_CONFIRM	DTS →CAS	Confirmation that a given C_PDU has been sent using the expedited delivery service
D_ EXPEDITED_UNIDATA_ REQUEST_REJECTED	DTS →CAS	Notification that a given C_PDU could not be sent using the expedited delivery service and the reason why it was rejected by Data transfer Sublayer.
D_EXPEDITED_UNIDATA_ INDICATION	DTS →CAS	Delivers a C_PDU that has been received using the expedited delivery service from a remote node, with additional indications of the transmission service given to the C_PDU, the source node address, and the destination node address (e.g., if a group address is the destination).
D_WARNING_RECEIVED	DTS →CAS	Notifies the Channel Access sublayer that a WARNING D_PDU has been received from a remote node by the Data Transfer Sublayer, with the reason for sending the warning message
D_WARNING_TRANSMITTED	DTS →CAS	Notifies the Channel Access sublayer that a WARNING D_PDU has been sent to a remote node by the Data Transfer Sublayer, with the reason for sending the warning message
D_CONNECTION_MADE	CAS →DTS	Notifies the Data Transfer Sublayer that a connection has been made with a given remote node.
D_CONNECTION_TERMINATED	CAS →DTS	Notifies the Data Transfer Sublayer that a connection has been terminated with a given remote node.
D_CONNECTION_LOST	DTS →CAS	Notifies the Channel Access Sublayer that a connection has been lost with a given remote node.

Note: DTS = Data Transfer Sublayer; CAS = Channel Access Sublayer; <from sublayer> —→ <to sublayer>

C.3.   Support for Edition 3 Interoperability

Edition 4 of STANAG 5066 introduces a number of new protocol elements to improve the DTS service. It is a key goal of Edition 4 to ensure that there is robust interoperability with   implementations following STANAG 5066 Edition 3 and earlier.

To achieve this an implementation **shall** determine if a peer implementation supports Edition 4 (or subsequent). If a peer only supports edition 3 or the status cannot be determined, an implementation **shall not** use any of the Edition 4 capabilities set out in this annex. This will ensure interoperability with Edition 3.

Peer Capability can be determined in three ways.

1. A priori knowledge.   The edition supported by a specific peer or by the whole network may be known.

2. Use of CAS-1 capability negotiation, as specified in the CAS-1 sub-layer. This is always used for ARQ data when ALE is not used. It **may** be used for ARQ data with ALE and **shall** be used if the peer capability is not known. This means that peer capability will always be known if ARQ data is being transferred.

3. The third mechanism is use of the Capability EOW defined in Section C.6.4. This enables Edition 4 peer capability to be explicitly determined with no ARQ data is being exchanged. It also facilitates capability update when a node is upgraded to Edition 4.

There are three types of Edition 4 capabilities that needs to be considered:

1.   Procedural options, which enhance performance, but **shall** only be used with Edition 4 (or subsequent) peers.

2. D_PDU Types summarized and categorized in Section C.4. There are three types of D_PDU.

    a. D_PDUs for use with all Editions.

    b. D_PDUs that **shall** only be used with Edition 4 and subsequent Editions.

    c. D_PDUs for use with Edition 3, that **may** be used in Edition 4 when both peers choose to.

3. EOWs are summarized and categorized in Section C.6**Error! Reference source not found.**.   There are three types of EOW:

    a. EOWs for use with all Editions.

    b. EOWs for use with Edition 3 only.

    c. EOWs for use with Edition 4 and subsequent Editions only.

Edition 4 capabilities **shall** only be used with peers known to support Edition 4.   This   ensures   interoperability with Edition 3 systems.

C.4.   Structure of Data Transfer Sublayer Protocol Data Units (D_PDUs)

In order to provide the data transfer services specified herein, the Data Transfer Sublayer **shall** [(1)] exchange protocol data units (D_PDUs) with its peer(s).

The Data Transfer Sublayer **shall** [(2)] use the D_PDU types displayed in Table C-2 to support the **Selective ARQ service** and **Non ARQ** service, including the several data transfer submodes defined herein.

Table C-2. D_PDU Types

D_PDU Frame Types	D_PDU Type #	Function	Protocol Type	Frame Type	Edition
ED3-DATA-ONLY	0	*ARQ data transfer*	SRQ	I	Ed3
ED3-ACK-ONLY	1	*Acknowledgement of ARQ data transfer*	SRQ	C	Ed3
ED3-DATA-ACK	2	*ARQ data transfer with acknowledgement*	SRQ	I+C	Ed3
ED3-RESET/WIN-RESYNC	3	*Reset/Re-synchronized peer protocol entities*	IRQ	C	Ed3
EXPEDITED-DATA-ONLY	4	*Expedited ARQ data transfer*	SRQ	I	All
EXPEDITED-ACK-ONLY	5	*Acknowledgement of expedited ARQ data transfer*	SRQ	C	All
MANAGEMENT	6	*Management message*	IRQ	C	All
NON-ARQ-DATA	7	*Non-ARQ data transfer*	NRQ	I	All
EXPEDITED-NON-ARQ-DATA	8	*Expedited non-ARQ data transfer*	NRQ	I	All
DATA-ONLY	9	*ARQ data transfer*	SRQ	I	Ed4
ACK-ONLY	10	*Acknowledgement of data transfer*	SRQ	C	Ed4
DATA-ACK	11	*ARQ data transfer with acknowledgement*	SRQ	I+C	Ed4
RESET/WIN-RESYNC	12	*Reset/Re-synchronized peer protocol entities*	IRQ	C	Ed4
EXTENSION	13	*Extension message*	-	-	Ed4
PADDING	14	*Header only*	-	C	Ed4
WARNING	15	*Unexpected or unrecognized D_PDU type*	-	C	All

C-Frame = Control Frame.   I-Frame = Information Frame.   I+C-Frame = Information + Control Frame.

Edition Support has the following options controlling STANAG 5066 Editions:

1. All: For use in all Editions; or
2. Ed4: For use in Edition 4 and subsequent editions only; or
3. Ed3: Primarily for use for Edition 3 interoperability but **may** be used between Edition 4 peers by mutual negotiated agreement.

All D_PDU types may be used to support single channel and full duplex transmission modes. There are basically three different types of D_PDUs, or frames, noted by the *Frame-Type* field in Table C-2:

1. C (Control) Frames; or
2. I (Information) Frames; or
3. A combined I+C Frame.

The *Protocol Type* field in Table C-2 indicates the type of data-transfer-service protocol with which the D_PDU frame **shall** [(3)] be used, as follows:

1. NRQ: No Repeat-Request.(i.e., Non-ARQ) Protocol; or

2. SRQ: Selective Repeat-Request Protocol; or
3. IRQ: Idle Repeat-Request Protocol

After sending I-frames, the NRQ protocol does not wait for an indication from the remote node as to whether or not the I-frames were correctly received. Multiple repetitions of I-frames can be transmitted in order to increase the likelihood of reception under poor channel conditions, in accordance with the requested service characteristics. Multiple transmissions of C frames is often desirable, as loss of acks can lead to unnecessary duplicate data transmission.

After sending I-frames, The Selective RQ protocol waits for an indication in the form of a selective acknowledgement from the remote node as to whether the I-frames were correctly received or not. The local node then either sends the next I-frame, if all the previous I-frames were correctly received, or retransmits copies of the previous I-frame that were not, in accordance with the requirements of Section C.7. The local node will retransmit copies of the previous I- frames if no indication is received after a predetermined time interval. A local node sending I-frames in full duplex mode also sends the indication in the form of a selective acknowledgement embedded in the I-frames as to whether the I-frames were correctly received or not.

After sending an I-frame, The Idle RQ protocol, also known as a stop and wait protocol must wait until it receives an acknowledgement from the remote node as to whether or not the I-frame was correctly received. The local node then either sends the next I-frame, if the previous I-frame was correctly received, or retransmits a copy of the previous I-frame if it was not. The local node will retransmit a copy of the previous I-frame if no indication is received after a predetermined time interval.

Different D_PDU frame types may be combined in a transmission, subject to limitations imposed by the state of the Data Transfer Sublayer protocol. These states of the Data Transfer Sublayer protocol and their associated requirements are given in Section C.7.1

C.4.1.   Generic simplified D_PDU structure

All D_PDU types that cannot carry segmented C_PDUs **shall** [1] be of the structure shown in Figure C-1 (a).

D_PDU types that can carry segmented C_PDUs shall (2) be structured according to Figure C-1 (b).

SYNC (16 bits)	HEADER (max 38 bytes)	HEADER CRC (16 bits)

Figure C-1 (a). Format for D_PDU C-Frame types (1, 3, 5, 6, 10, 12, 13, 14 and 15)

SYNC (16 bits)	HEADER (max 38 bytes)	HEADER CRC (16 bits)	SEGMENTED N_PDU (max 1023 bytes)	CRC (32 bits)

Figure C-1 (b). Format for D_PDU I and I+C Frame types (0, 2, 4, 7, 8, 9, 10, 11 and 13)

C.4.2. Generic detailed D_PDU structure

The detailed structure of the generic D_PDU C-Frame **shall** [1] be as shown in Figure C-2 (a) or Figure C-2(b)

The D_PDU types 1, 3, 5, 6, 12, 14 and 15and 15 **shall** [2] use only the C-Frame structure defined in Figure C-2 (a). D_PDU types 10 and 13 **may** use this structure.
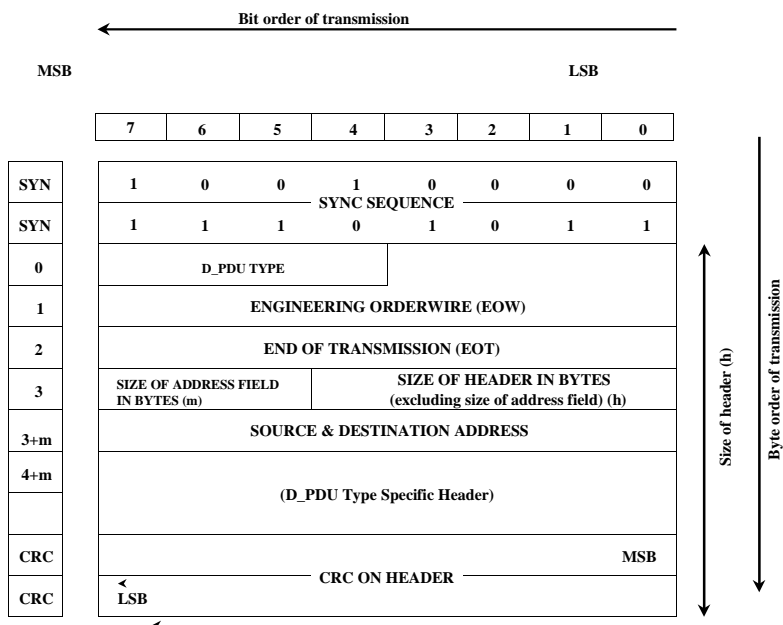

Figure C-2 (a). Generic D_PDU C-Frame Structure.

The D_PDU types 0, 2, 4, 7, 8, 9, 11 **shall** [(3)] use the generic D_PDU I and I+C Frame structure defined in Figure C-2 (b). D_PDU types 10 and 13 **may** use this structure.

Bit order of transmission
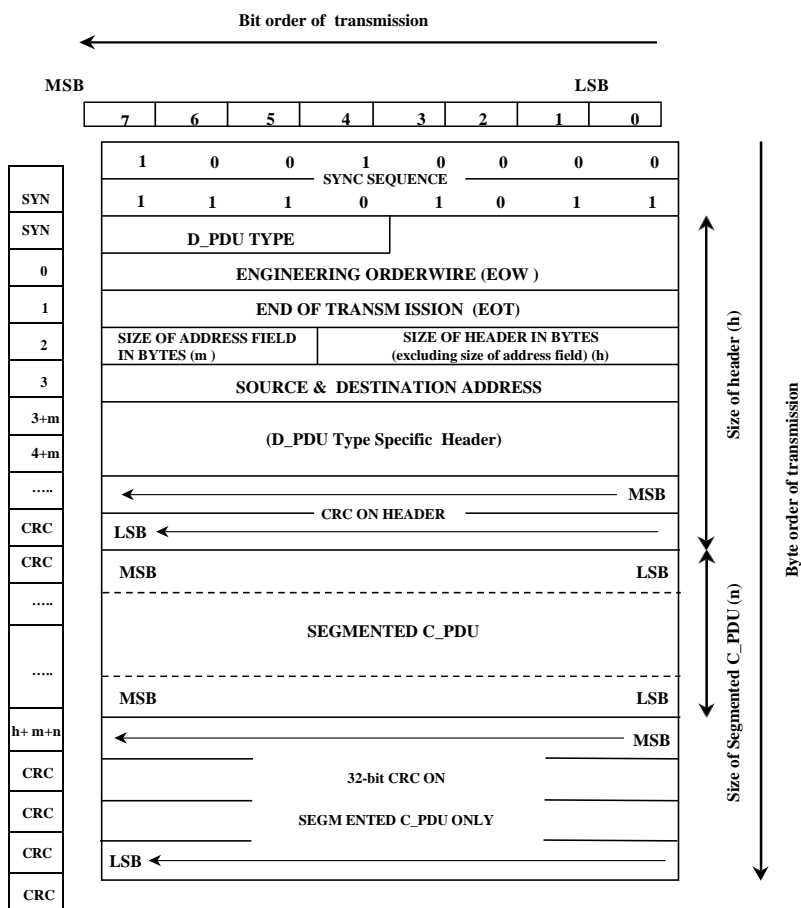
Figure C-2 (b). Generic D_PDU I and I+C Frame Structure.

All D_PDUs, regardless of type, **shall** [(4)] begin with the same 16-bit synchronisation (SYNC) sequence.

The 16-bit sequence **shall** [(5)] be the 16-bit Maury-Styles (0xEB90) sequence shown below, with the least significant bit (LSB) transmitted first:

(MSB) 1 1 1 0 1 0 1 1 1 0 0 1 0 0 0 0 (LSB)

16-bit Maury-Styles synchronisation sequence

The first 4 bytes of all D_PDU headers **shall** [(6)] contain the same fields:
1. a 4 bit *D_PDU Type* field that **shall** [(7)] identify the type of D_PDU;
2. a 12-bit field that **shall** [(8)] contain an engineering order wire (*EOW*) message;
3. an 8-bit field that **shall** [(9)] contain the end of transmission (*EOT*) information; and
4. a one byte field that **shall** [(10)] contain both a *Size-of-the-Address* field

(3 bits) and a *Size-of-the-Header* (5 bits) field.

The next 1 to 7 bytes of every header, as specified in the *Size-of-the-Address* field, **shall** [11] contain source and destination address information for the D_PDU.

The D_PDU *Type-Specific-Header-Part* field **shall** [12] be as specified below in this STANAG, for each of the D_PDU types.

The last two bytes of every header **shall** [13] contain the Cyclic Redundancy Check (CRC) calculated in accordance with Section C.4.2.8.

The bits in any field in a D_PDU that is specified as NOT USED **shall** [14] contain the value zero (0).

C.4.2.1. D_PDU type

The D_PDU types **shall** [1] be as defined in Table C-2 and the D_PDU figures below.

The value of the D_PDU type number **shall** [2] be used to indicate the D_PDU type. The four bits available allow for 16 D_PDU types. All the possible values are assigned in this edition of STANAG 5066.

C.4.2.2. Engineering Orderwire (EOW)

The 12 bit EOW field **shall** [1] carry Management messages for the Engineering Orderwire (EOW). EOW messages may not be explicitly acknowledged although the D_PDU of which they are a part may be. EOW messages can be explicitly acknowledged when they are contained in the MANAGEMENT Type 6 D_PDU through which Management-level acknowledgement services are provided in the Data Transfer Sublayer.

Figure C-3 (a) shows the generic 12-bit EOW structure. The first 4 bits of the EOW **shall** [2] contain the EOW-type field, which identifies the type of EOW message. The remaining 8-bits **shall** [3] contain the EOW- type-specific EOW data.

The various EOW messages including their definitions and extensions are specified further in Section C.6.

Figure C-3 (b) shows how the EOW message is mapped into the generic D_PDU header.

Bit mapping order

| MSB											LSB
11	10	9	8	7	6	5	4	3	2	1	0
EO W TYPE				EO W specific data							

Figure C-3 (a). Generic EOW message.

Figure C-3 (b). EOW mapping convention in D_PDU header.

C.4.2.3. End Of Transmission (EOT)

The 8-bit EOT field **shall** [1] provide an approximation of the time remaining in the current transmission interval specified by the transmitting node. This information is provided for the timing of exchanges between peer nodes to minimize collisions and the link turnaround time (time between the end of a transmission by one node and the start of a transmission by another node).

The number in this field **shall** [2] be a binary number expressing the number of half (1/2) second intervals remaining in the current transmission from the beginning of the current D_PDU including sync bytes.

In some modes of transmission, as described in Section C.7.5, the EOT is set to zero. When the EOT is set to zero, no information is provided as to when the transmission ends.

C.4.2.4. Size of Address Field

The Size-of-Address Field **shall** [1] specify the number of bytes in which the source and destination address are encoded (Note: this value is denoted by the integer value "m" in Figure C-2(a) and Figure C-2(b)). The address field may be from one (1) to seven (7) bytes in length, with the source and destination address of equal length.

Since the D_PDU header must be made up of an integer number of bytes, addresses **shall** [2] be available in 4-bit increments of size: 4 bits (or 0.5 bytes), 1 byte, 1.5 bytes, 2 bytes, 2.5 bytes, 3 bytes, and 3.5 bytes.

C.4.2.5. Size of Header Field

The Size-of-Header field **shall** [1] specify the number of bytes in which the D_PDU is encoded. (Note: this value is denoted by the integer value "h", Figure C-2(a) and Figure C-2(b)), and its value includes the sizes of the following fields and elements:

        C.4.2.5.1.1. D_PDU Type
        C.4.2.5.1.2. EOW
        C.4.2.5.1.3. EOT
        C.4.2.5.1.4. Size of Address field
        C.4.2.5.1.5. Size of Header field
        C.4.2.5.1.6. D_PDU-Type-specific header
        C.4.2.5.1.7. CRC field

The value of the Size-of-Header field **shall** [2] not include the size of the source and destination address field.

C.4.2.6. Source & Destination Address

Each D_PDU transmitted by a node **shall** [(1)] contain the source and destination address. Half of the bits are assigned to the source and the other half to the destination.

The first half **shall** [(2)] be the destination address and the second half **shall** [(3)] be the source address as displayed nominally in Figure C-4(a) (which assumes an odd-number as the address-field size) or Figure C- 4(b) (which assumes an even-number as the address-field size).


Figure C-4 (a). Address mapping convention in D_PDU header, assuming address-field size is odd

.
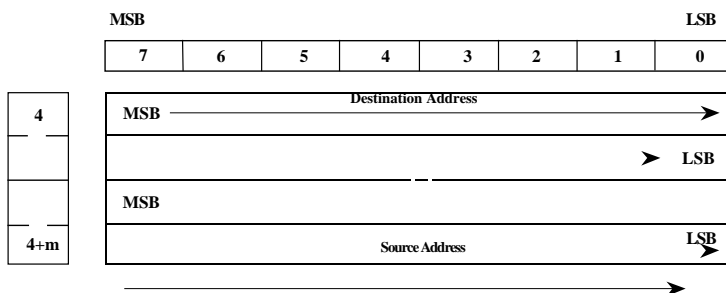
Figure C-4(b). Address mapping convention in D_PDU header, assuming address-field size is even.

Addresses **shall** [(4)] be in the form of a binary number. With 7 bytes available for each of the user and the destination, the smallest possible address field is 4 bits, the largest possible is 3.5 bytes, or 28 bits.

A decimal number **shall** [(5)] represent each byte or fractional byte of an address, and the binary equivalent **shall** [(6)] be mapped into the corresponding byte. Consequently, the decimal representation for node addresses (i.e., unicast addresses) lies in the range
[0.0.0.0 ... 15.255.255.255] (this **may** be referred to as 'dotted-decimal address format').

For DPDUs that support group addressing (e.g; THE Type 7/8 Non-ARQ DPDUs), the group-address flag **should** be considered the 29th-bit in the binary representation of the address. Consequently, the decimal representation for group addresses (i.e., multicast addresses) lies in the range [16.0.0.0 ... 31.255.255.255].

The broadcast (all-stations) address in "dotted-decimal" address format is the group address 31.255.255.255.

Any fractional-byte elements in the address **shall** [7] be mapped into the first (leftmost) non-zero number in the in the decimal representation of the address. The remaining numbers in the decimal representation of the address **shall** [8] refer to byte-sized elements in the address field. [Note: As an example, if 3.5 bytes are used, the address would be expressed as w.x.y.z, where w can be any value from 0 to 15, and x, y and z can be any value from 0 to 255. The value w will represent the most significant bits and z the least significant bits of the address.]

The address bits **shall** [9] be mapped into the address field by placing the MSB of the address into the MSB of the first byte of the address field, and the LSB into the LSB of the last byte of the field, in accordance with Figure C-4(a), for addresses with length of 0.5, 1.5, 2.5, or 3.5 bytes, and Figure C-4(b), for addresses with length of 1, 2, or 3 bytes.

When a field spans more than one octet, the order of the bit values within each octet **shall** [10] decrease progressively as the octet number increases.

The lowest bit number associated with the field represents the lowest-order value. Leading address bytes which are zero may be dropped from the address, consistent that the requirement that the source and destination address subfields must be of equal length. Trailing address bytes that are zero **shall** [11] be sent.

C.4.2.7. D_PDU Type-Specific Header

The bytes immediately following the address field **shall** [1] encode the D_PDU Type-Specific header, as specified in the corresponding section below from Sections C.4.4 through C.4.15.

C.4.2.8. Cyclic Redundancy Check (CRC)

A two byte header-only CRC is used. Because the D_PDU header is generally shorter than the data, errors are more likely in the data part of a D_PDU than the header. Protecting the header with its own CRC allows the possibility to detect and use uncorrupted header information even if the data part of a D_PDU contains errors. For this reason, the added overhead of the CRC-on-header field was deemed warranted in the design of STANAG 5066.

The two-bytes following the D_PDU Type-Specific header **shall** [1] contain a 16-bit Cyclic Redundancy Check (CRC) field.

The header CRC error-check field **shall** [2] be calculated using the following polynomial: $x^{16} + x^{15} + x^{12} + x^{11} + x^8 + x^6 + x^3 + 1$, or in hexadecimal format 0x19949, using the shift-register method shown by the figures in Appendix I of CCITT Recommendation V.41 (or equivalent method in software; an example is given below).
[Note: This polynomial is not the same as that shown in the figure in V.41, Appendix I; the polynomial was chosen to provide a lower probability of undetected error (i.e., better performance) than that given by the polynomial specified in V.41. The polynomial specified here was analyzed and reported by Wolf in a paper in the 1988 IEEE Conference on Military Communications (MILCOM- paper 15.2). Note too that the taps in the figure are shown reversed from the order in which they occur in the polynomial defined in V.41. This reversal is accommodated by the remaining requirements specified below.]

When calculating the header CRC field, the shift registers **shall** [3] be initially set to all (0) zeros.

The header CRC **shall** [4] be calculated over all bits in the header, excluding the Maury-Styles synchronisation

sequence, and including the following fields and elements:

- D_PDU Type
- EOW
- EOT
- Size of Address field
- Size of Header field
- Source and Destination Address
- D_PDU-Type-Specific header

A node **shall** [5] process the information contained in a header with a valid CRC, regardless of the result of the CRC error check over any segmented C_PDU that may be a part of the D_PDU.

The CRC bits **shall** [6] be mapped (see Figure C-5) into the CRC octets by placing the MSB of the CRC into the LSB of the first byte of the CRC field, and the LSB of the CRC into the MSB of the last byte of the CRC field. This will result in the MSB of the most significant byte of the CRC being sent first, followed by the remaining bits in descending order, which is consistent with the order of transmission for CRC bits, as specified in Recommendation V.42, section 8.1.2.3.


Figure C-5. CRC mapping convention in D_PDU header.

The following C code can be used to calculate the CRC value using the specified polynomial.

```
/* Polynomial x^16 + x^15 + x^12 + x^11 + x^8 + x^6 + x^3 + 1        */
unsigned short CRC_16_S5066(unsigned char DATA, unsigned short CRC)
{
unsigned char i, bit;
for (i=0x01; i; i<<=1)
{
bit = (((CRC & 0x0001) ? 1:0)^((DATA&i) ? 1:0));
CRC>>=1;
if (bit) CRC^=0x9299;          /* polynomial representation, bit */
}                              /* -reversed (read right-to-left), and */
                               /* with the initial x^16 term implied. */

return (CRC);
}


/*   example of usage: */
#define NUM_OCTETS;       /* number of octets in the message data */
unsigned char        message[NUM_OCTETS];
unsigned short       CRC_result;
unsigned int j;


CRC_result = 0x0000;
for (j=0; j < NUM_OCTETS; j++){
       CRC_result = CRC_16_S5066(message[j], CRC_result);
       }
/* CRC_result contains final CRC value when loop is completed */
```

NOTE: *This code calculates the CRC bytes in the proper order for transmission as defined above; no bit reversal is required with this code.*

Code Example C-1. Calculation of the CRC-16-S5066 for D_PDU Headers.

The function CRC_16_5066 in Code Example C-1 may be used to calculate the CRC for a data sequence of octets, and is called for each successive octet in the sequence. The function accepts as its first argument an octet in the data sequence, and as its second argument, the value of the CRC calculated for the previous octet in the sequence. The function returns the value of the CRC. When called for the first octet in the data sequence, the value of the CRC must be initialized to zero as required.

The result of this code for the D_PDU described below is 0x1E5F. The least significant byte of the computed CRC should be transmitted before the most significant byte; as noted, the algorithm already performs the requisite bit-level reversal. The receive processing should extract the CRC bytes and re-assemble into the correct order.

```
        D_PDU used in above CRC calculation:
        D_PDU Type:        Warning
        EOW:               Type 0 (i.e., all zeros)
        EOT:               all zeros
        Address size:      2
```

Destination address:	0.0.0.5 (leading zeros not encoded)
Source address:	0.0.0.100 (leading zeros not encoded)
Received D_PDU type:	0
Reason warning sent:	2
CRC:	0x1E5F

The full D_PDU is then (including sync bytes and CRC-field, properly bit-reversed and in the order they should be transmitted):

0x90, 0xEB, 0xF0, 0x00, 0x00, 0x47, 0x05, 0x64, 0x02, 0x5F, 0x1E

C.4.2.9. Segmented C_PDU

For the I and I+C D_PDUs types, the octets of the segmented C_PDUs **shall** [1] be transmitted in ascending numerical order, following the two-byte CRC on the D_PDU header.

Within an octet, the LSB **shall** [2] be the first bit to be transmitted as shown in Figure C-6.
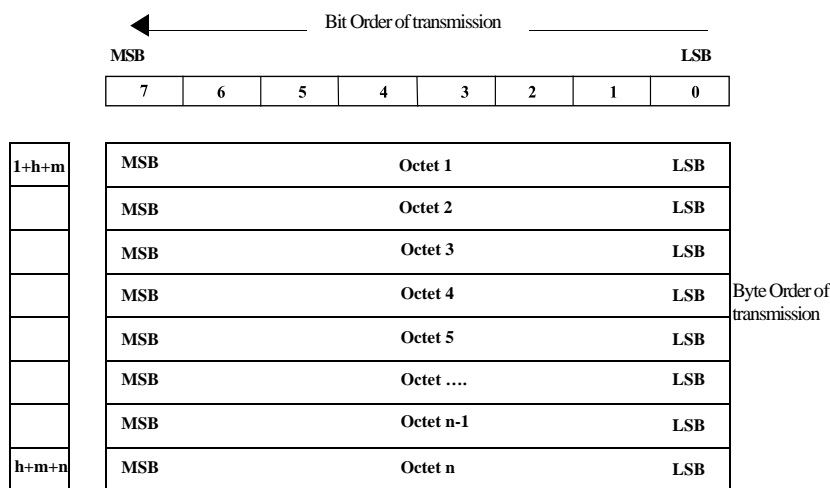

Figure C-6. Segmented C_PDU mapping convention in D_PDU structure.

C.4.2.10. Size of Segmented C_PDU

The SIZE OF SEGMENTED C_PDU field **shall** [1] be used only with DATA D_PDUs that are I or I+C frame types, i.e, that have a Segmented C_PDU field as shown in Figure C-2(b). (Note: The value of the SIZE OF SEGMENTED C_PDU field is denoted by the integer value "n" in the Figure). This field actually is contained within the D_PDU Type-Specific Header part, but since it is common to all I and I+C D_PDUs the format is defined here.

The bit-value of the SIZE OF SEGMENTED C_PDU **shall** [2] be encoded as a ten-bit field as indicated by Figure C-7.

Bi t mapping order

MSB

LSB

| 9	8	7	6	5	4	3	2	1	0
SI ZE OF SEGMENTED N_PDU IN BYTES									

Figure C-7. SIZE OF SEGMENTED C_PDU Field.

The value in the SIZE OF SEGMENTED C_PDU field **shall** [3] not include the two-bytes for the CRC following the Segmented C_PDU. The Segmented C_PDU field can hold a maximum of 1023 bytes from the segmented C_PDU.

The SIZE OF SEGMENTED C_PDU **shall** [4] be mapped into consecutive bytes of the D_PDU as indicated in

Figure C-8, in the byte locations specified for the applicable D_PDU.

MSB   LSB

7	6	5	4	3	2	1	0

4+m		MSB ⟶
5+m		⟶ LSB

Figure C-8. Segmented C_PDU Size mapping convention in all applicable D_PDU header structures.

C.4.2.11.   CRC-ON-SEGMENTED-C_PDU Field

The last four bytes of any I or I+C D_PDU **shall** [1] contain a 32-bit Cyclic Redundancy Check (CRC) field.

The CRC **shall** [2] be applied and computed on the contents of the Segmented C_PDU using the following polynomial [see footnote 2 below]:

$$x^{32} + x^{27} + x^{25} + x^{23} + x^{21} + x^{18} + x^{17} + x^{16} + x^{13} + x^{10} + x^8 + x^7 + x^6 + x^3 + x^2 + x + 1$$

or, in hexadecimal notation: 0x10AA725CF, using the shift-register method similar to that shown by the figures in Appendix I of CCITT Recommendation V.41, but using a longer shift register and appropriate changes to the tap configuration corresponding to the polynomial specified. Equivalent implementations in software may be used to compute the 32-bit CRC (an example is offered below).

NOTE: This polynomial is constructed with a methodology similar to that used to construct the 16-bit polynomial on the header, and for the same reason. Using this methodology, the polynomial for the 32-bit CRC on the segmented C_PDU is the generator polynomial for a two-error-correcting BCH code (65535,32), with a maximum information block length of 65503 bits. Only the error-detection properties of the code are used.

When calculating the header CRC field, the shift registers **shall** [4] be initially set to all (0) zeros.

The following C code can be used to calculate the CRC value using the specified polynomial.
```
/* Polynomial:
    x^32 + x^27 + x^25 + x^23 + x^21 + x^18 + x^17 + x^16 + x^13 + x^10 + x^8 + x^7 + x^6 + x^3 + x^2 + x + 1 */
        unsigned int CRC_32_S5066(unsigned char DATA, unsigned int CRC)
        {
        unsigned char i, bit;

        for (i=0x01; i; i<<=1)
        {
        bit = (((CRC & 0x0001) ? 1:0)^((DATA&i) ? 1:0));
        CRC>>=1;
        if (bit) CRC^=0xF3A4E550;            /* polynomial representation, bit */
        }                                    /* -reversed (read right-to-left), and */
                                             /* with the initial x^32 term implied. */
        return (CRC);
        }


        /*  example of usage: */
        #define NUM_OCTETS;           /* number of octets in the message data */
        unsigned char message[NUM_OCTETS];
        unsigned int    CRC_result;
        unsigned int    j;

        CRC_result = 0x00000000;
        for (j=0; j < NUM_OCTETS; j++){
                CRC_result = CRC_32_S5066(message[j], CRC_result);
                }
        /* CRC_result contains final CRC value when loop is completed */
```

NOTE: *This code calculates the CRC bytes in the proper order for transmission as defined above;*
no bit reversal is required with this code.

Code Example C-2. Calculation of the CRC-32-S5066 on Segmented C_PDUs.

The function CRC_32_S5066 in Code Example C-2 may be used to calculate the CRC for a data sequence of octets, and is called for each successive octet in the sequence. It is initialized and used in the same manner as the example for the STANAG 5066 16-bit CRC.

When applied to this short C_PDU message data sequence:

message[] = {0xF0, 0x00, 0x00, 0x47, 0x05, 0x64, 0x02},

the result is the 32-bit value (in hexadecimal format):   0xF4178F95

Just as for the 16-bit CRC computation, the algorithm that calculates the 32-bit CRC performs the bit-level reversal in place, and the resultant value must be transmitted least-significant byte first, in order to most-significant byte last. The full message data, with 32-bit CRC appended in proper position, is the following

sequence:
    0xF0, 0x00, 0x00, 0x47, 0x05, 0x64, 0x02, 0x95, 0x8F, 0x17, 0xF4

C.4.3.   CHOICE OF FRAME SEQUENCE NUMBER LENGTH FOR ARQ D_PDUs

STANAG 5066 Ed3 used an eight bit frame sequence number for ARQ data. This choice leads to window exhaustion and significant performance impact at wideband HF speeds and at faster narrowband speeds. Edition 4 introduces a set of four PDUs to support regular ARQ data with a 16 bit frame sequence number, which addresses the performance issue. For three of the PDUs, the longer frame sequence number is the only change. ACK-ONLY has some additional changes to optimize performance.

These PDUs are described in pairs in Sections C.4.4 - C.4.7. The Edition 3 compatible PDUs are prefixed with "ED3-".   These "ED3-" PDUs **shall** be used for regular ARQ data exchange with Edition 3 and earlier peers.

It is anticipated that the new D_PDUs will usually be used for regular ARQ communication between Edition 4 and subsequent peers and this choice is the default. Use of the Edition 3 D_PDUs **may** be negotiated using CAS-1 link setup with agreement of both peers. This has potential to gain about 0.5% throughput improvements in links which will only operate at the slowest narrowband HF speeds.

It is anticipated that Edition 5 will make support for these Edition 3 D_PDUs optional and it is possible that support will be removed in future editions.

C.4.4. DATA-ONLY D_PDUs (Simplex data transfer)
C.4.4.1. ED3-DATA-ONLY (Type 0) D_PDU
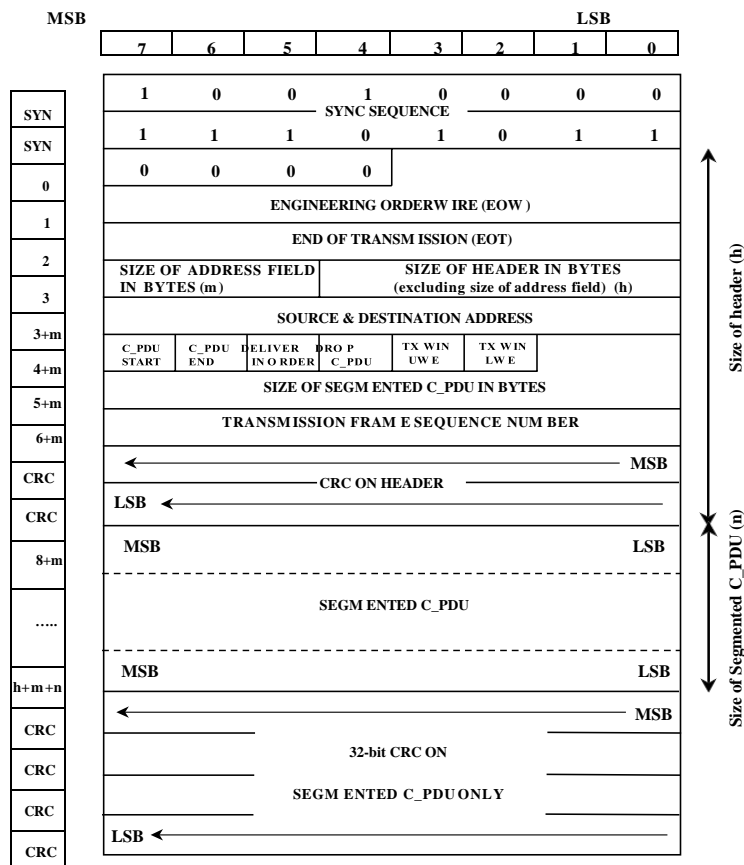
Figure C-9. Frame format for ED3-DATA-ONLY D_PDU Type 0

The encoding of the ED3-DATA-ONLY D_PDU is show in Figure C-9.

C.4.4.2. DATA-ONLY (Type 9) D_PDU

	MSB 7	6	5	4	3	2	1	LSB 0
0	DPDU Type = 9				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m	C_PDU START	C_PDU END	DELIVER IN ORDER	DROP C_PDU	TX WIN UWE	TX WIN LWE		
5+m	Segment Size (n)							
6+m*	Transmit Frame Sequence Number						MSB	
7+m*	LSB	(increased 8 to 16 bits)						
CRC	CRC on Header							
CRC								
h+m	Segmented C_PDU							
h+m+n-1								
CRC	CRC on Data							
CRC								
CRC								
CRC								

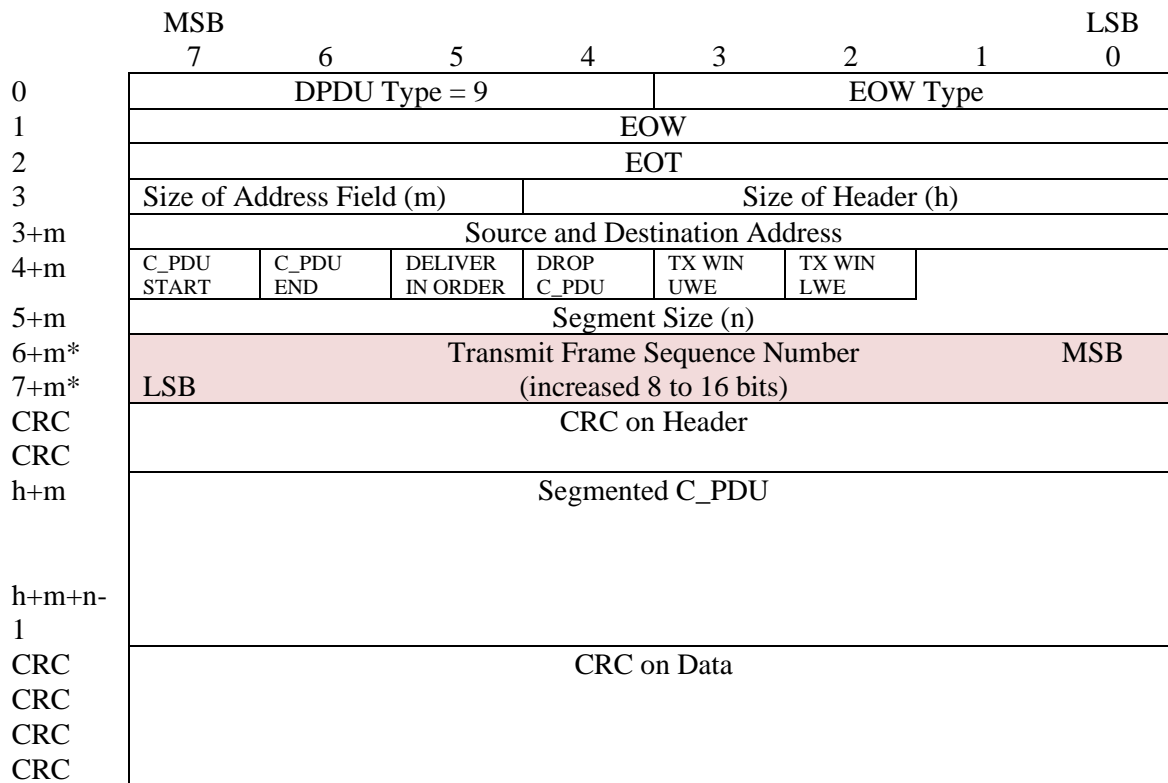**Figure C-10. Frame format for DATA-ONLY D_PDU Type 9**

The encoding of the DATA-ONLY D_PDU is shown in Figure C-10. The differences relative to ED3-DATA-ONLY D_PDU are highlighted and numbering changes marked with "*".

C.4.4.3. DATA-ONLY Fields

The term "DATA-ONLY D_PDU" in this section is used to refer to both Type 0 and Type 9 D_PDUs.

The DATA-ONLY D_PDU **shall** [1] be used to send segmented C_PDUs when the transmitting node needs an explicit confirmation the data was received.

The DATA-ONLY D_PDU **shall** [2] be used in conjunction with a basic selective automatic repeat request type of protocol.

A Data Transfer Sublayer entity that receives a DATA-ONLY D_PDU **shall** [3] transmit an ACK-ONLY (TYPE 1) D_PDU or a DATA-ACK (TYPE 2) D_PDU as acknowledgement, where the type of D_PDU sent depends on whether or not it has C_PDUs of its own to send to the source of

the DATA-ONLY D_PDU.

The DATA-ONLY D_PDU **shall** [(4)] contain the following fields within its D_PDU Type-Specific part, mapped and encoded in accordance with Figure C-9, Figure C-10, and the paragraphs below:

- C_PDU START
- C_PDU END
- DELIVER IN ORDER
- DROP C_PDU
- TX WIN UWE
- TX WIN LWE
- SIZE OF SEGMENTED C_PDU
- TRANSMIT SEQUENCE NUMBER

The C_PDU START flag **shall** [(5)] be set to indicate the start of a newly segmented C_PDU; the C_PDU segment contained within this D_PDU is the first segment of the C_PDU, in accordance with the C_PDU-segmentation process described in section C.5

The C_PDU END flag **shall** [(6)] be set to indicate the end of a segmented C_PDU; when a D_PDU is received with the C_PDU END flag set it indicates the last D_PDU that was segmented from the C_PDU. Depending on the status of the DELIVER IN ORDER flag, the link layer will assemble and deliver the C_PDU, if all D_PDUs between and including the C_PDU START and C_PDU END are received completely error free. The re-assembly process of D_PDUs into C_PDUs is described in section C.5

If the DELIVER IN ORDER flag is set on the D_PDUs composing a C_PDU, the C_PDU **shall** [(7)] be delivered to the upper layer when both the following conditions are met:

1) The C_PDU is complete.
2) All C_PDUs received previously that also had the DELIVER IN ORDER flag set have been delivered.

If the DELIVER IN ORDER flag is cleared on the D_PDUs composing a C_PDU, the C_PDU **shall** [(8)] be delivered to the upper layer when the following condition is met:

3) The C_PDU is complete and error free.

The DROP C_PDU flag is used when the TTL of a C_PDU expires before transmission of the C_PDU is complete. It is necessary to send and acknowledge this D_PDU in order to maintain window synchronization. A D_PDU with the DROP C_PDU flag set **shall** be acknowledged.

No segmented C_PDU data **shall** be sent if the DROP C_PDU flag is set as the data is being discarded and the SIZE OF SEGMENTED C_PDU field **shall** [(10)] be zero in this case.

When the DROP C_PDU flag is set by the D_PDU source, the receiving Data Transfer Sublayer **shall** [(9)] discard the contents of the segmented C_PDU field of the current D_PDU and all other previously received segments of the C_PDU of which the current D_PDU is a part.

The TX WIN UWE flag **shall** [11] be set when the TRANSMIT FRAME SEQUENCE NUMBER for the current D_PDU is equal to the Transmit Window Upper Edge (TX UWE) of the transmit-flow-control window.

Similarly, the TX WIN LWE flag **shall** [12] be set when the TRANSMIT FRAME SEQUENCE NUMBER for the current D_PDU is equal to the Transmit Lower Window Edge (LWE) of the transmit flow control window.

The SIZE OF SEGMENTED C_PDU field **shall** [13] be encoded as specified in Section C.4.2.10.

The TRANSMIT FRAME SEQUENCE NUMBER field **shall** [14] contain the sequence number of the current D_PDU.

The value of the TRANSMIT FRAME SEQUENCE NUMBER field **shall** [15] be a unique integer assigned to the D_PDU during the segmentation of the C_PDU, and will not be released for reuse with another D_PDU until the receiving node has acknowledged the D_PDU. This integer **shall** be modulo 256 for Type 0 D_PDU and modulo 65,536 for type 9.

Values for the TRANSMIT FRAME SEQUENCE NUMBER field **shall** [16] be assigned in an ascending order of appropriate modulo during the segmentation of the C_PDU.

The SEGMENTED C_PDU field **shall** [17] immediately follow the D_PDU header as depicted in Figure C-7. Segmented C_PDUs **shall** [18] be mapped according to the specification of Section C.4.2.9 .

C.4.5. ACK-ONLY (TYPE 1 or TYPE 10) D_PDU (Acknowledgement of type 0, 2 data transfer)
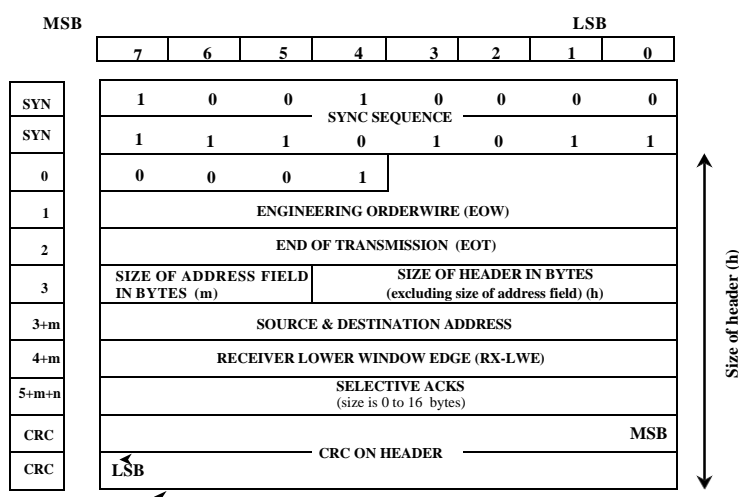
C.4.5.1. ED3-ACK-ONLY (Type 1) D_PDU

Figure C-11. Frame format for ED3-ACK-ONLY D_PDU Type 1

The encoding of the ED3-DATA-ONLY D_PDU is shown in Figure C-11.

C.4.5.2. ACK-ONLY (Type 10) D_PDU

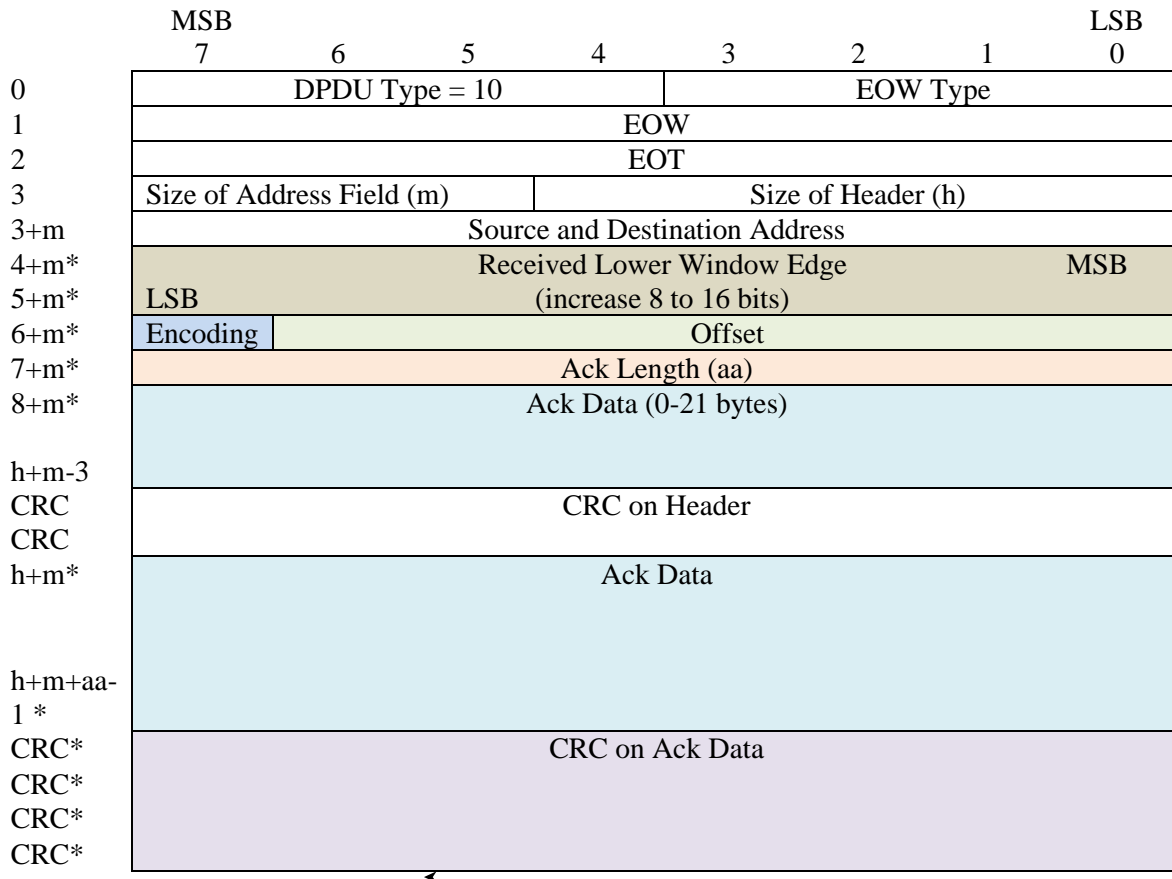|	MSB							LSB
	7	6	5	4	3	2	1	0
0	DPDU Type = 10				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m*	Received Lower Window Edge						MSB	
5+m*	LSB	(increase 8 to 16 bits)						
6+m*	Encoding	Offset						
7+m*	Ack Length (aa)							
8+m*	Ack Data (0-21 bytes)							
h+m-3								
CRC	CRC on Header							
CRC								
h+m*	Ack Data							
h+m+aa-1 *								
CRC*	CRC on Ack Data							
CRC*								
CRC*								
CRC*								

Figure C-12. Frame format for ACK-ONLY D_PDU Type 10

The encoding of the ACK-ONLY D_PDU is shown in Figure C-12. The differences relative to ED3-ACK-ONLY D_PDU are highlighted and numbering changes marked with "*".

C.4.5.3. ACK-ONLY D_PDU Common Encoding

The term "ACK-ONLY D_PDU" in this section is used to refer to both Type 1 and Type 10 D_PDUs.

The ACK-ONLY D_PDU **shall** [1] be used to selectively acknowledge received DATA-ONLY or DATA-ACK D_PDUs when the receiving Data Transfer Sublayer has no segmented C_PDUs of its own to send or (Edition 4 only) when necessary acknowledgement information cannot be encoded in DATA-ACK D_PDU.

The ACK-ONLY D_PDU **shall** [2] contain the following fields within its D_PDU Type-Specific part, mapped and encoded in accordance with Figure C-11 or Figure C-12 and the paragraphs below:

    ☐   RECEIVE LOWER WINDOW EDGE (LWE)

The value of the RECEIVE LOWER WINDOW EDGE (RX LWE) field **shall** [3] equal the D_PDU sequence number of the RX LWE pointer associated with the node's receive ARQ flow-control window**.** This integer **shall** be modulo 256 for Type 1 D_PDU and modulo 65,536 for type 10.

C.4.5.4. ED3-ACK-ONLY D_PDU Encoding for Type 1 only

The ED3-ACK-ONLY D_PDU **shall** [2] contain the following fields within its D_PDU Type-Specific part, mapped and encoded in accordance with Figure C-11 and the paragraphs below:

    ☐   SELECTIVE ACKS

The SELECTIVE ACKS field can have a dynamic length of 0 to 16 bytes, and **shall** [4] contain a bit-mapped representation of the status of all received D_PDUs with sequence numbers from the LWE to and including the UWE pointers of the receive flow-control window. The size of the SELECTIVE ACK field can be determined from knowledge of the structure of the ED3-ACK-ONLY D_PDU and the value of the Size of Header field, and is not provided explicitly in the ED3-ACK-ONLY D_PDU.

A set (1) bit within the SELECTIVE ACKS field **shall** [5] indicate a positive acknowledgement (ACK), i.e., that the D_PDU with the corresponding Frame Sequence Number was received correctly.

Only D_PDU frames with a correct segmented C_PDU CRC **shall** [6] be acknowledged positively even if the header CRC is correct, except that frames with the DROP C_PDU flag set **shall** [7] be acknowledged positively regardless of the results of the CRC check on the segmented C_PDU.

A cleared (0) bit within the SELECTIVE ACKS field **shall** [8] indicate a negative acknowledgement (NACK), i.e., that the D_PDU with the corresponding Frame Sequence Number was received incorrectly, or not at all.

The construction of the SELECTIVE ACK field and the mapping of D_PDU frame-sequence numbers to bits within the SELECTIVE ACK field **shall** [9] be in accordance with Figure C-13, Figure C-14 and the paragraphs below.
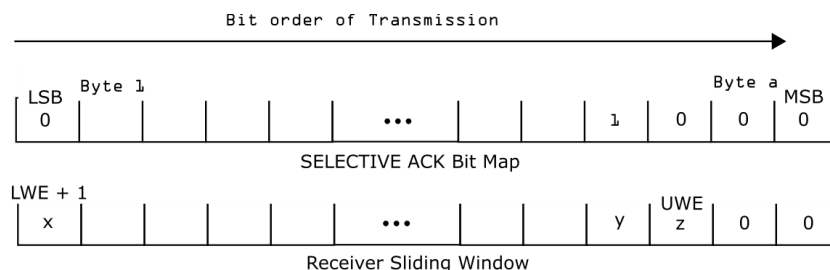
Bit order of Transmission

| | Byte 1 | | | | ... | | | | | | Byte a | MSB |
| LSB 0 | | | | | | | 1 | 0 | 0 | 0 | | |

SELECTIVE ACK Bit Map

| LWE + 1 | | | | | ... | | | | UWE | | |
| x | | | | | | | y | z | 0 | 0 | |

Receiver Sliding Window

Figure C-13. Constructing the SELECTIVE ACK Field
(Note: the bits that are set (1) and cleared (0) are representing example bits)

MSB							LSB
7	6	5	4	3	2	1	0

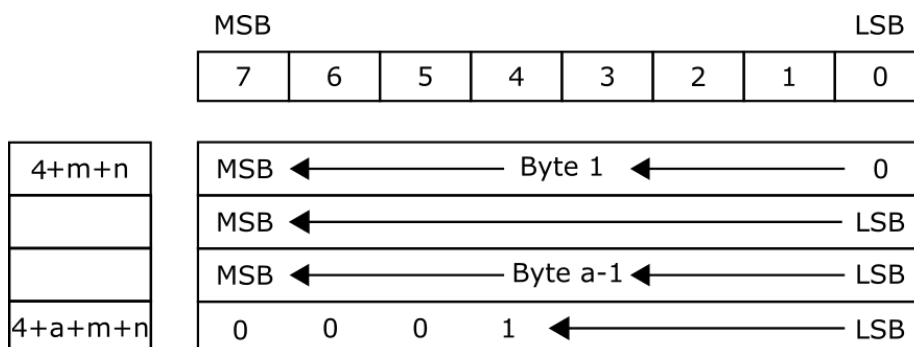| 4+m+n | MSB ← Byte 1 ← 0 |
	MSB ← LSB
	MSB ← Byte a-1 ← LSB
4+a+m+n	0 0 0 1 ← LSB

Figure C-14. SELECTIVE ACK mapping convention
(Note: the example bits that were shown in Figure C-13 can been seen in this mapping example)

The LSB of the first byte of the SELECTIVE ACK field **shall** [10] correspond to the D_PDU whose frame sequence number is equal to 1 + the value in the RX LWE field of this ED3-ACK-ONLY D_PDU [the bit corresponding to the RX LWE is always zero, by definition, and is therefore not sent].

Each subsequent bit in the SELECTIVE ACK field **shall** [11] represent the frame sequence number of a subsequent D_PDU in the receive flow-control sliding window, in ascending order of the respective frame-sequence numbers without omission of any values.

The bit corresponding to the Receive Upper Window Edge (RX UWE) **shall** [12] be in the last byte of the SELECTIVE ACK field.

If the bit representing the RX UWE is not the MSB of the last byte, the remaining bits in the byte (until and including the MSB) **shall** [13] be set to 0 as padding. No further SELECTIVE ACK bytes **shall** [14] be transmitted after such bits are required.

C.4.5.5. ACK-ONLY D_PDU Encoding for Type 10 only

The Ack Data in the ACK-ONLY D_PDU encodes acknowledgements of transmitted D_PDU frames. In order to efficiently encode acknowledgements when the window size is large, two different encodings are defined. Acknowledgements **may** be encoded in multiple ACK-ONLY D_PDUs in order to keep the size of each D_PDU smaller. It is recommended to keep ACK_ONLY D_PDU size small in order to reduce risk of D_PDU loss during transmission. It is also recommended to transmit ACK_ONLY D_PDUs multiple times, because loss of acknowledgement information can lead to delays and unnecessary retransmission.

When a set (1 or more) of ACK-ONLY D_PDUs is sent all of the received DATA PDUs **shall** be acknowledged. The number and encoding of these D_PDUs is an implementation choice.

Only D_PDU frames with a correct segmented C_PDU CRC **shall** [6] be acknowledged positively even if CRC is correct, except that frames with the DROP C_PDU flag set **shall** [7] be acknowledged positively the header regardless of the results of the CRC check on the segmented C_PDU.

If Ack Data can be encoded in 21 bytes or less the D_PDU encoding of Figure C-1(a) is used. All of the Ack Data is encoded in the D_PDU header.

If Ack Data is larger, the encoding with two CRCs of Figure C-1(b) is used. In this case no Ack data **shall** be encoded in the header. All of the Ack Data **shall** be encoded in the main D_PDU after the header.

The maximum window size can lead to up to 32,768 packets to be acknowledged, which would need up to 4096 bytes with the standard encoding mechanism. It is generally be desirable to keep ACK-ONLY D_PDUs smaller than this, in order to reduce risk of D_PDU loss. The offset mechanism provides a means to do this, by enabling reporting acknowledgements by referencing D_PDUs relative to an offset of the Lower Window Edge. Multiple ACK-ONLY D_PDUs **may** be sent with different offsets, in order to acknowledge all D_PDUs received. The Offset is encoded as 7 bits in the ACK-ONLY D_PDU header. The Offset value is multiplied by 256 to determine the size of the offset in bits (where each bit represents a PDU that is being acknowledged). This enables the acknowledgement data to be reduced, so that the Ack Data **may** be reduced so that it will never exceed 32 bytes.

The size in bytes of Ack Data used is specified in Ack Length encoded in the ACK-ONLY D_PDU header. It **may** be up to 255 bytes.

Two options are provided to encode Ack Data, with the choice controlled by the Encoding Bit in the ACK-ONLY D_PDU header. A sender **may** choose to use either or both encodings. A recommended approach is for the sender to evaluate both encodings and then to use the more compact one. A receiver **shall** support both encodings.

The first mechanism is the one defined in Section C.4.5.4. The Encoding bit is set to 0 for use of this encoding. This encoding sets a bit is set to 1 to indicate that a Data PDU has been accepted. The data starts with reference to the first PDU after the Offset. Bit encoding follows the rules of Section C.4.5.4. It is safe to add trailing zeros to pad to an exact number of bytes.

The second is a simple Run Length encoding mechanism described in this section. It is designed as a byte-aligned format that can efficiently encode long runs of 1s or 0s, which are anticipated to be likely when the window size is large.

Run Length Encoding starts with the standard Ack encoding defined in Section C.4.5.4 and generates a different encoding. In some common situations, this will be more compact. It is recommended to use the more compact encoding. The encoding is illustrated in Figure C-15 and Figure C-16.

	MSB							LSB
	7	6	5	4	3	2	1	0
0	Value	Extended=0	MSB		Length			LSB

Figure C-15. Run Length Encoding: Single Byte Encoding

|	MSB							LSB
	7	6	5	4	3	2	1	0
0	Value	Extended=1	MSB		Length			
1								LSB

Figure C-16. Run Length Encoding: Two Byte Encoding

The first bit (Value) is set to 0 or 1 and shows the value that is being repeated. If the Extended bit is set to 0, the encoding uses one byte and Length is specified in 6 bits. If the Extended bit is set to 1, the encoding uses two bytes and Length is specified in 14 bits. The Length indicates the number of repeats of the value, so if length is zero the value occurs once and is not repeated.

These encodings are repeated through Ack Data.

C.4.6.   DATA-ACK D_PDU
C.4.6.1. ED3-DATA-ACK (Type 2) D_PDU

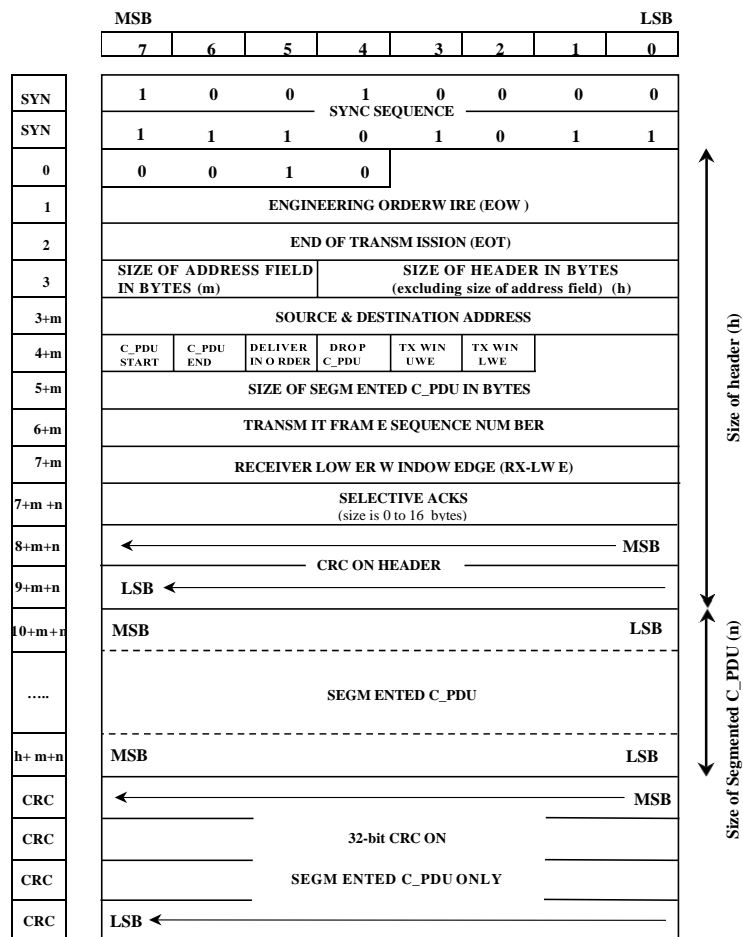**Figure C-17. Frame format for ED3-DATA-ACK D_PDU Type 2**

The encoding of the ED3-DATA-ACK D_PDU is shown in

Figure C-17.

C.4.6.2. DATA-ACK (Type 11) D_PDU

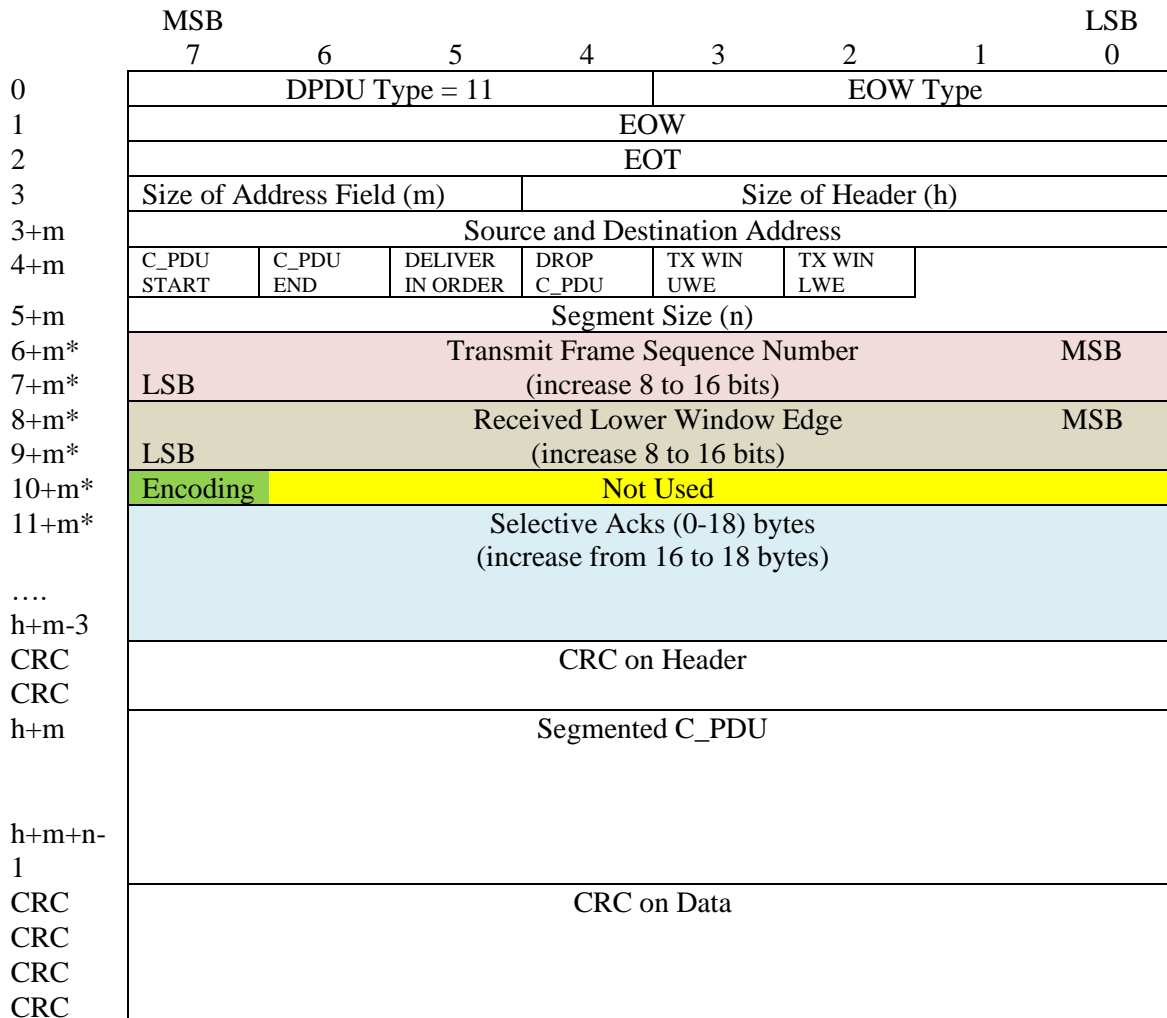|	MSB 7	6	5	4	3	2	1	LSB 0
0	DPDU Type = 11				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m	C_PDU START	C_PDU END	DELIVER IN ORDER	DROP C_PDU	TX WIN UWE	TX WIN LWE		
5+m	Segment Size (n)							
6+m*	Transmit Frame Sequence Number						MSB	
7+m*	LSB	(increase 8 to 16 bits)						
8+m*	Received Lower Window Edge						MSB	
9+m*	LSB	(increase 8 to 16 bits)						
10+m*	Encoding	Not Used						
11+m*	Selective Acks (0-18) bytes (increase from 16 to 18 bytes)							
….								
h+m-3								
CRC CRC	CRC on Header							
h+m	Segmented C_PDU							
h+m+n-1								
CRC CRC CRC CRC	CRC on Data							

Figure C-18. Frame format for DATA-ACK D_PDU Type 11

The encoding of the DATA-ACK D_PDU is shown in Figure C-18. The differences relative to ED3-DATA-ACK D_PDU are highlighted and numbering changes marked with "*".

C.4.6.3. DATA-ACK D_PDU Common Encoding

The term "DATA-ACK D_PDU" in this section is used to refer to both Type 2 and Type 11 D_PDUs.

DATA-ACK D_PDU **shall** be used when data and acknowledgements are being transmitted, rather than DATA-ONLY plus ACK-ONLY D_PDUs. If only one acknowledgement is being sent, use of DATA_ACK D_PDU is preferred as overhead is reduced. All of the acknowledgement information is in the header, so the acknowledgement information can be valid when data is corrupted in transfer. If acknowledgements are repeated, it is recommended to use ACK-ONLY

D_PDUs to do this.

The DATA-ACK (TYPE 2) D_PDU is a part of a basic selective automatic repeat request type of protocol. The DATA-ACK D_PDU **shall** [1] be used to send segmented C_PDUs when the transmitting node needs an explicit confirmation the data was received and has received D_PDUs to selectively acknowledge.

A Data Transfer Sublayer entity that receives a DATA-ACK D_PDU **shall** [2] transmit an ACK-ONLY D_PDU or a DATA-ACK D_PDU as acknowledgement, where the type of D_PDU sent depends on whether or not it has C_PDUs of its own to send to the source of the DATA-ACK D_PDU.

The DATA-ACK D_PDU is a combination of the DATA-ONLY and ACK-ONLY D_PDU. All of the field specifications from the DATA-ONLY and ACK-ONLY D_PDUs apply to this D_PDU. The DATA-ACK D_PDU **shall** [4] contain the following fields within its D_PDU Type-Specific part, mapped and encoded in accordance with

Figure C-17 and the referenced paragraphs:

- C_PDU START – **shall** be as specified in Section C.4.4.3 for the DATA-ONLY D_DPU;
- C_PDU END– **shall** be as specified in Section C.4.4.3.C.4.4 for the DATA-ONLY D_DPU;
- DELIVER IN ORDER– **shall** be as specified in Section C.4.4.3 for the DATA-ONLY D_DPU;
- DROP C_PDU– **shall** be as specified in Section C.4.4.3 for the DATA-ONLY D_DPU;
- TX WIN UWE– **shall** be as specified in Section C.4.4.3C.4.4 for the DATA-ONLY D_DPU;
- TX WIN LWE– **shall** be as specified in Section C.4.4.3 for the DATA-ONLY D_DPU;
- SIZE OF SEGMENTED C_PDU– **shall** be as specified in Section C.4.4.3C.4.4 for the DATA-ONLY D_DPU;
- TRANSMIT SEQUENCE NUMBER– **shall** be as specified in Section C.4.4.3 for the DATA-ONLY D_DPU;
- RECEIVE LOWER WINDOW EDGE (RX LWE) – **shall** be as specified in Section C.4.5.3 for the ACK- ONLY D_DPU;

C.4.6.4. Encoding Selective Acks

For ED3-DATA-ACK (Type 2), SELECTIVE ACKS **shall** be as specified in Section C.4.5.4 for the ED3-ACK-ONLY D_DPU.

For DATA-ACK(Type 2), a subset of the rules of Section C.4.5.5 shall be followed as described here. The Encoding bit controls the choice of encoding in the same manner as in Section C.4.5.5.

Selective Acks **shall** be encoded in the manner of the Ack Data encoding set out in Section C.4.5.5, controlled by the encoding choice. Selective Acks is limited to 18 bytes, which may prevent acknowledgement of D_PDUs a long way from the receiving window edge. In this situation, ACK-ONLY D_PDU **shall** be used to acknowledge these D_PDUs.

C.4.7. RESET/WIN-RESYNC D_PDU (Reset/Re-synchronise peer protocol entities)

C.4.7.1. ED3-RESET/WIN-RESYNC (TYPE 3) D_PDU

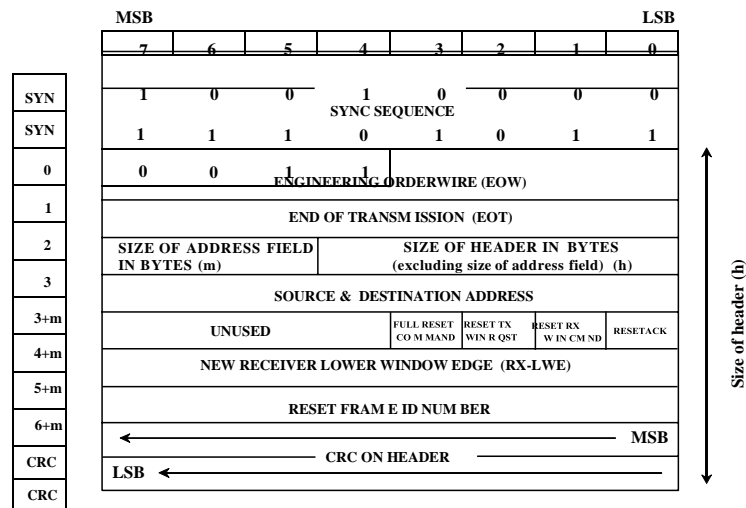**Figure C-19. Frame format for ED3-RESET/WIN-RESYNC D_PDU Type 3**

The encoding of the ED3-RESET/WIN_RESYNC D_PDU is show in Figure C-19.

C.4.7.2. RESET/WIN-RESYNC (TYPE 12) D_PDU

	MSB 7	6	5	4	3	2	1	LSB 0
0	\multicolumn DPDU Type = 12				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m	Unused			FRC	RTWQ	RRWC		RSA
5+m*	New Receiver Lower Window Edge							MSB
6+m*	LSB	(increase 8 to 16 bits)						
7+m	Reset Frame ID Number							
CRC	CRC on Header							
CRC								

Figure C-20. Frame format for RESET/WIN-RESYNC D_PDU Type 12

The encoding of the RESET/WIN-RESYNC D_PDU is shown in Figure C-10Figure C-20. The differences relative to ED3- RESET/WIN-RESYNC D_PDU are highlighted and numbering changes marked with "*".

C.4.7.3. RESET/WIN-RESYNC Fields

The term "RESET/WIN-RESYNC D_PDU" in this section is used to refer to both Type 3 and Type 12 D_PDUs.

The RESET/WIN-RESYNC D_PDU **shall** [1] be used to control the re-synchronisation or re-initialisation of the selective-repeat ARQ protocol operating on the link between the source and destination nodes.

Reset and resynchronization operations **shall** [2] be performed with respect to the transmit lower-window edge (TX LWE) and the receive lower-window edge (RX LWE) for the flow-control sliding windows at the sending node and receiving node, as specified in Section C.7.2.

The RESET/WIN-RESYNC D_PDU **shall** [3] use a basic stop and wait type of protocol (denoted as the IRQ protocol elsewhere in this STANAG) in which the reception of this D_PDU **shall** [4] result in the transmission of an acknowledgement D_PDU by the receiving node. For the IRQ protocol used with the RESET/WIN-RESYNC D_PDU, the RESET/WIN-RESYNC D_PDU is used for both data and acknowledgement, as specified below.

Transmission of D_PDUs supporting the regular-data service, i.e., of DATA, ACK, and DATA-ACK D_PDUs, **shall** [5] be suspended pending completion of any stop-and-wait protocol using the RESET/WIN-RESYNC D_PDUs. [Note: The ARQ windowing variables will be reset to zero or changed as a result of the use of the REST/WIN-RESYNC D_PDUs, and suspension of normal data transmission until these variables have been reset is required, lest the variables controlling the selective-repeat ARQ protocol become even further corrupted than the error-state that presumably triggered the reset or resynchronization protocol.]

The RESET/WIN RESYNC D_PDU **shall** [6] contain the following fields within its D_PDU Type-Specific part, mapped and encoded in accordance with Figure C-19 and the paragraphs below:

- FULL RESET COMMAND
- RESET TX WIN RQST
- RESET RX WIN CMND
- RESET ACK
- NEW RECEIVE LOWER WINDOW EDGE (LWE)
- RESET FRAME ID NUMBER

The FULL RESET COMMAND flag **shall** [7] be set equal to one (1) to force a full reset of the ARQ machines at the transmitter and receiver to their initial values as specified in Sections C.7.2 and C.7.3.

A RESET/WIN-RESYNC D_PDU with the RESET TX WIN RQST flag set equal to one (1) **shall** [8] be used to request a resynchronisation of the TX-LWE and RX-LWE pointers used for DATA in the transmit and receive nodes.

A node that receives a RESET/WIN-RESYNC D_PDU with the RESET TX WIN RQST flag set equal to one **shall** [9] respond by forcing resynchronization of the windows using a RESET/WIN RESYNC D_DPU and the RESET RX WIN CMND flag, as specified below.

A RESET/WIN-RESYNC D_PDU with the RESET RX WIN CMND flag set equal to one (1) **shall**
be used to force a resynchronisation of the TX-LWE pointer at the sending node and the RX-LWE pointer at the receiving node.

A node that sends a RESET/WIN-RESYNC D_PDU with the RESET RX WIN CMND flag set equal to one **shall** [11] proceed as follows:
- The NEW RECEIVE LWE field **shall** [12] be set equal to the value of the sending node's TX-LWE.
- The sending node **shall** [13] wait for a RESET/WIN-RESYNC D_PDU with the RESET ACK flag set equal to one as an acknowledgement that the resynchronization has been performed.

A node that receives a RESET/WIN RESYNC D_PDU with the RESET RX WIN CMND flag set equal to one **shall** [14] proceed as follows:
- The value of the node's TX LWE **shall** [15] be set equal to the value of the NEW RECEIVE LWE field in the RESET/WIN RESYNC D_PDU that was received;
- The node **shall** [16] send a RESET/WIN-RESYNC D_PDU with the RESET ACK flag set equal to one as an acknowledgement that the resynchronization has been performed.

The RESET ACK flag **shall** [17] be set equal to one (1) to indicate an acknowledgement of the most recently received RESET/WIN-RESYNC D_PDU.

A node may use a RESET/WIN-RESYNC acknowledgement transmission, i.e., a RESET/WIN-RESYNC D_PDU with the RESET ACK flag set equal to one (1), to request/force a reset or resynchronization of its own ARQ flow control variables, e.g., if the link is supporting two way data communication and the ARQ machines for both directions of data-flow must be reset or resynchronized.

The NEW RECEIVE LWE field specifies the value of the new receiver ARQ RX-LWE, as noted above, and

shall [18] be valid only when the value of the RESET RX WIN CMND flag equals one (1). The value of the NEW RECEIVE LWE field **shall** [19] be ignored in any other situation.

The Data Transfer Sublayer **shall** [20] use the RESET FRAME ID NUMBER field to determine if a given RESET/WIN-RESYNC D_PDU received is a copy of one already received.

The value of the RESET FRAME ID NUMBER field **shall** [21] be a unique integer of appropriate modulo assigned in ascending order to RESET/WIN RESYNC D_PDUs, and will not be released for reuse with another D_PDU until the D_PDU to which it was assigned has been acknowledged.
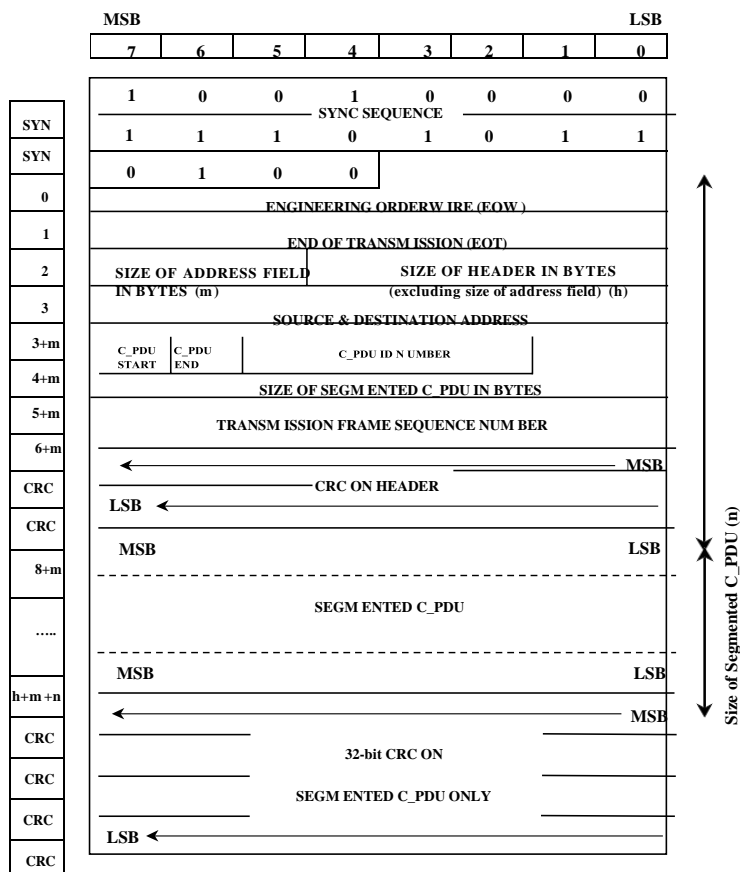
C.4.8.   EXPEDITED-DATA-ONLY (TYPE 4) D_PDU

Figure C-21. Frame format for EXPEDITED DATA ONLY D_PDU Type 4

The EXPEDITED-DATA-ONLY (TYPE 4) D_PDU **shall** [1] be used to send segmented C_PDUs that require Expedited Delivery Service when the transmitting node needs an explicit confirmation the data was received.

A Data Transfer Sublayer entity that receives EXPEDITED-DATA-ONLY (TYPE 4) D_PDU **shall** send an EXPEDITED-ACK-ONLY (TYPE 5) D_PDU as a selective acknowledgement of all EXPEDITED-DATA-ONLY (TYPE 4) D_PDUs received from the source node.

The EXPEDITED-DATA-ONLY D_PDU is similar in structure to the ED3-DATA-ONLY D_PDU. The EXPEDITED-DATA-ONLY D_PDU **shall** [3] contain the following fields within its D_PDU Type-Specific part, mapped and encoded in accordance with Figure C-21 and the paragraphs noted:

- C_PDU START – **shall** [4] be as specified for the DATA-ONLY D_PDU in Section C.4.4
- C_PDU END – **shall** [5] be as specified for the DATA-ONLY D_PDU in Section C.4.4
- C_PDU ID NUMBER – **shall** [6] be as specified in the paragraphs below.
- SIZE OF SEGMENTED C_PDU– **shall** [7] be as specified in Section C.4.2.10 for all D_PDUs that have a Segmented C_PDU field.
- TRANSMIT SEQUENCE NUMBER– **shall** [8] be as specified for the DATA-ONLY D_PDU in Section C.4.4, with additional requirements as noted below.

The C_PDU ID NUMBER field **shall** [6] specify the C_PDU of which this Expedited D_PDU is a part. The value of the C_PDU ID NUMBER field **shall** [7] be an integer (modulo 16) assigned in an ascending modulo 16 order to the C_PDU, and **shall** [8] not be released for reuse with another C_PDU until the entire C_PDU has been acknowledged.

As noted above, the TRANSMIT SEQUENCE NUMBER field in the EXPEDITED-DATA-ONLY D_PDU is defined and used in the same manner as that specified for the ED3-DATA-ONLY D_PDU. However, the EXPEDITED-DATA-ONLY D_PDUs **shall** [9] be assigned frame numbers from a frame sequence number pool (0, 1 ...255) that is reserved exclusively for the transmission of EXPEDITED-DATA-ONLY and EXPEDITED-ACK-ONLY D_PDUs. The FRAME SEQUENCE NUMBER is used, in this D_PDU, to sequence the D_PDUs that make up a C_PDU receiving expedited delivery service.
(Note: the further implication of this requirement is that there are independent state machines and flow-control windows [or different states and sets of variables within a single state-machine] for the Expedited and Regular delivery services in the Data Transfer Sublayer).

The SEGMENTED C_PDU field is a field that is attached to the header structure defined in Figure C-21. The segmented PDU **shall** [10] immediately following the D_PDU header. Segmented C_PDUs **shall** [11] be mapped according to the convention described in C.3.2.9.

The processing of EXPEDITED D_PDUs in the EXPEDITED DATA state **shall** [11] differ from the processing of DATA-ONLY or DATA-ACK D_PDUs in the DATA state in the following ways:

- Data (i.e, C_PDUs) using the Expedited Delivery Service **shall** [12] be transferred using EXPEDITED-DATA-ONLY and EXPEDITED-ACK-ONLY D_PDUs. If two way data communication is required, EXPEDITED-DATA-ONLY and EXPEDITED-ACK-ONLY D_PDUs may be placed together in a transmission interval.

- C_PDUs requiring Expedited Delivery Service and the associated EXPEDITED D_PDUs **shall** [13] not be queued for processing within the Data Transfer Sublayer behind D_PDUs containing non-expedited data (i.e., DATA-ONLY or DATA-ACK D_PDUs).

C.4.9. EXPEDITED-ACK-ONLY (TYPE 5) D_PDU (Acknowledgement of expedited data transfer)

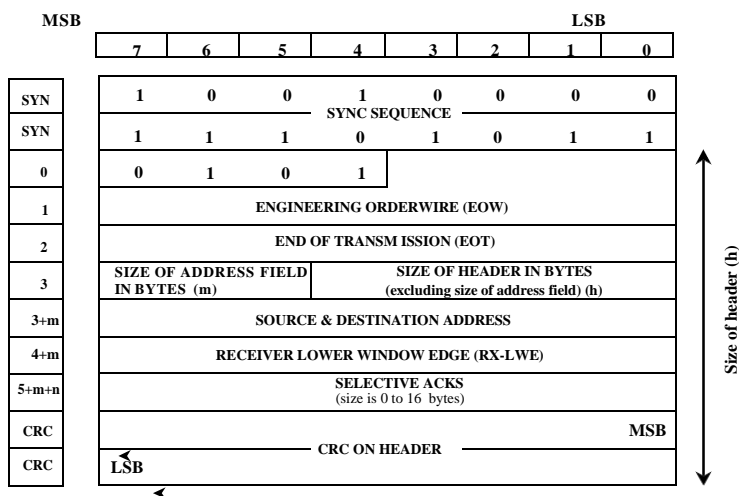
Figure C-22. Frame format for EXPEDITEDACK-ONLY D_PDU Type 5

The EXPEDITED-ACK-ONLY (TYPE 5) D_PDU **shall** [1] be used to selectively acknowledge received EXPEDITED-DATA-ONLY D_PDUs.

The EXPEDITED-ACK-ONLY (TYPE 5) D_PDU type **shall** [2] have the same format as the ED3-ACK-ONLY (TYPE 1) D_PDU, differing only in the value of the D_PDU Type field in byte 0, as specified in Figure C-22.

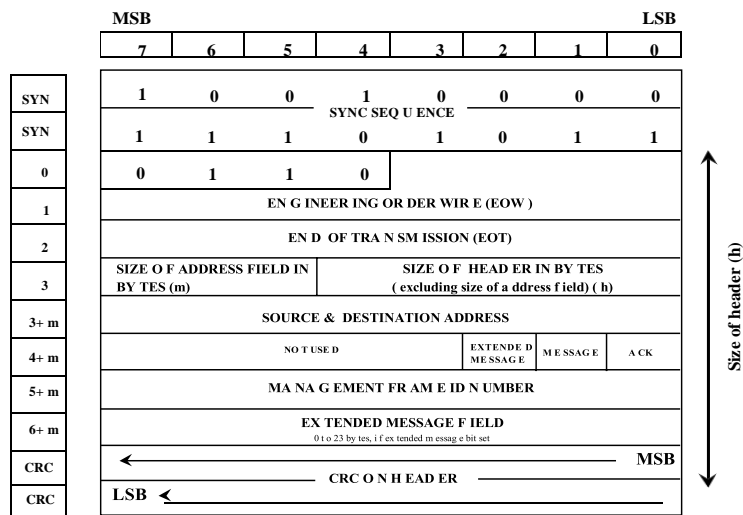### C.4.10. MANAGEMENT (TYPE 6) D_PDU (Management message transfer)

Figure C-23. Header format for MANAGEMENT D_PDU TYPE 6

The MANAGEMENT (TYPE 6) D_PDU **shall** [(1)] be used to send EOW Messages or
Management Protocol Data Units (M_PDUs) when the transmitting node needs an
explicit Acknowledgement that they were received.
MANAGEMENT D_PDU transmission bypasses other D_PDUs, so that they can be
delivered as quickly as possible.   Repeating transmission of MANAGEMENT
D_PDUs is recommended.

A Data Transfer Sublayer entity **shall** [(2)] acknowledge receipt of a MANAGEMENT (TYPE 6)
D_PDU by sending a MANAGEMENT (TYPE 6) D_PDU with the ACK flag set to the value
one (1).

The processing and transmission of MANAGEMENT (TYPE 6) D_PDUs **shall** [(3)] take
precedence over and bypass all other pending D_PDU types in the Data Transfer Sublayer.

The exchange of MANAGEMENT D_PDUs is regulated by a stop-and-wait protocol, i.e., there
shall [(4)] be only one unacknowledged MANAGEMENT D_PDU at any time.

The MANAGEMENT D_PDU **shall** [(5)] contain the following fields within its D_PDU Type-
Specific part, mapped and encoded in accordance with Figure C-23 and the paragraphs below:

- EXTENDED MESSAGE Flag
- VALID MESSAGE
- ACK
- MANAGEMENT FRAME ID NUMBER
- EXTENDED MANAGEMENT MESSAGE

The VALID MESSAGE field **shall** [(6)] be set to the value one (1) if the EOW field of the D_PDU
contains a valid Management message or the initial segment of a valid Management message
that is continued in the EXTENDED MANAGEMENT MESSAGE field.

The VALID MESSAGE field **shall** [7] be set to the value zero (0) if the EOW field contains an Engineering Orderwire Message for which an acknowledgement message is not required. If the VALID MESSAGE field is set to zero, the MANAGEMENT D_PDU **shall** [8] be used only to acknowledge receipt of another MANAGEMENT D_PDU.

The EXTENDED MESSAGE Flag **shall** [9] be set to the value one (1) if the D_PDU contains a non-zero, non-null EXTENDED MANAGEMENT MESSAGE field.

If the EXTENDED MESSAGE Flag is set to the value zero (0), the EXTENDED MANAGEMENT MESSAGE field **shall** [10] not be present in the MANAGEMENT D_PDU.

The MANAGEMENT FRAME ID NUMBER field **shall** [11] contain an integer in the range [0,255] with which MANAGEMENT D_PDUs **shall** [12] be identified.

The Data Transfer Sublayer **shall** [13] maintain variables to manage the frame ID numbers associated with this D_PDU:
- the TX MANAGEMENT FRAME ID NUMBER **shall** [14] maintain the value of the Frame ID Number for MANAGEMENT D_PDUs that are transmitted;
- the RX MANAGEMENT FRAME ID NUMBER **shall** [15] maintain the value of the Frame ID Number for the most recently received MANAGEMENT D_PDUs.

On initialisation (such as a new connection), a node's Data Transfer Sublayer **shall** [16] set its current TX MANAGEMENT FRAME ID NUMBER to zero and **shall** [17] set its current RX MANAGEMENT FRAME ID NUMBER to an out-of-range value (i.e., a value greater than 255).

The current value of the TX MANAGEMENT FRAME ID NUMBER **shall** [18] be placed in the appropriate field of each unique MANAGEMENT D_PDU transmitted. The current value of the TX MANAGEMENT FRAME ID NUMBER **shall** [19] be incremented by one, modulo 256, after each use, unless transmission of repeated copies of the MANAGEMENT D_PDU are specified for its use.

Management D_PDUs that have been repeated **shall** [20] have the same MANAGEMENT FRAME ID NUMBER.

The Data Transfer Sublayer **shall** [21] compare the MANAGEMENT FRAME ID NUMBER of received MANAGEMENT D_PDUs to the current RX MANAGEMENT FRAME ID NUMBER, and process them as follows:

- if the MANAGEMENT FRAME ID NUMBER in the received D_PDU differs from the current RX MANAGEMENT FRAME ID NUMBER value, the D_PDU **shall** [22] be treated as a new D_PDU, and the Data Transfer Sublayer **shall** [23] set the current RX MANAGEMENT FRAME ID NUMBER value equal to the value of the received MANAGEMENT FRAME ID NUMBER.

- if the value in the received D_PDU is equal to the current RX MANAGEMENT FRAME ID NUMBER value, the node **shall** [24] assume that the frame is a repetition of a MANAGEMENT D_PDU that has already been received, and the value of the current RX MANAGEMENT FRAME ID NUMBER **shall** be left unchanged.

There **shall** [26] be a one-to-one correspondence between MANAGEMENT messages and MANAGEMENT

D_PDUs; that is, each message is placed into a separate D_PDU (which may be repeated a number of times).

The 12-bit EOW section of the D_PDU **shall** [27] carry the EOW (non-extended) MANAGEMENT message, as specified in Section C.6.

The EXTENDED MANAGEMENT MESSAGE field may be used to transmit other implementation-specific messages that are beyond the scope of this STANAG. When the EXTENDED MESSAGE field is present and in use, the EXTENDED MESSAGE Flag **shall** [28] be set to the value one (1).

C.4.11. NON-ARQ-DATA (TYPE 7) D_PDU (Non-ARQ data transfer)
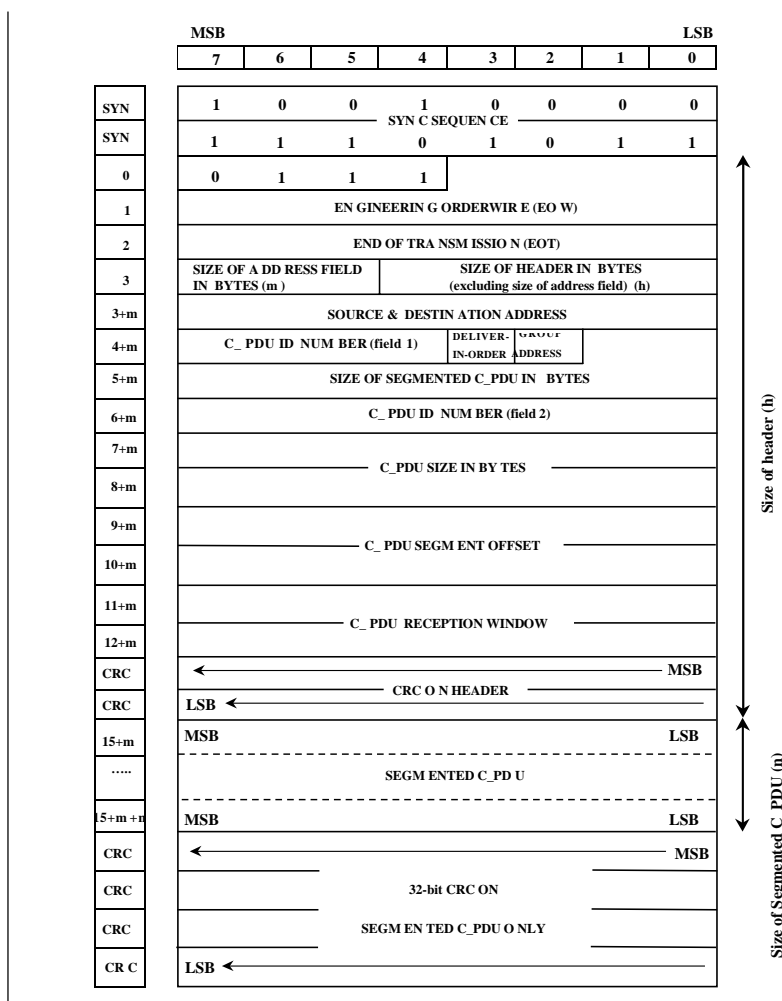

Figure C-24. Frame format for NON-ARQ-DATA D_PDU TYPE 7

The NON-ARQ-DATA (TYPE 7) D_PDU **shall** [1] be used to send segmented C_PDUs when the transmitting node needs no explicit confirmation the data was received.

The NON-ARQ-DATA D_PDU **shall** [2] contain the following fields within its D_PDU Type-Specific part, mapped and encoded in accordance with

Figure C-24 and the paragraphs below:

- C_PDU ID NUMBER (field 1)
- DELIVER IN ORDER
- GROUP ADDRESS
- SIZE OF SEGMENTED C_PDU
- C_PDU ID NUMBER (field 2)
- C_PDU SIZE
- C_PDU SEGMENT OFFSET
- C_PDU RECEPTION WINDOW

The C_PDU ID NUMBER field **shall** [3] identify the C_PDU to which the C_PDU segment encapsulated by the NON-ARQ DATA D_PDU belongs.

The value encoded in the C_PDU ID NUMBER field **shall** [4] be a unique integer (modulo 4096) identifier assigned in an ascending order (also modulo 4096) to the C_PDU during its segmentation and encapsulation into D_PDUs.

The value encoded in the C_PDU ID NUMBER field **shall** [5] not be released for reuse and assignment to another C_PDU until the time specified in the C_PDU RECEPTION WINDOW expires, as noted below.

The C_PDU ID NUMBER space (i.e, the set of ID numbers in the range [0..4095]) for NON-ARQ-DATA (TYPE 7) D_PDUs **shall** [6] be different than the similarly-defined number space for EXPEDITED-NON- ARQ-DATA (TYPE 8) D_PDUs.

The value of the C_PDU ID NUMBER **shall** [7] be encoded in a 12 bit field as specified in Figure C-25.

Bit mapping order

| MSB											LSB
11	10	9	8	7	6	5	4	3	2	1	0
N_PDU ID NUMBER											

Figure C-25. C_PDU ID NUMBER Field

The value of the C_PDU ID NUMBER **shall** [8] be mapped into the NON-ARQ DATA D_PDU into two split fields as follows, and as depicted in Figure C-26:

- The four-most-significant bits of the value of the C_PDU ID NUMBER **shall** [9] be mapped into C_PDU ID NUMBER (field 1);
- The eight least-significant bits of the value of the C_PDU ID NUMBER **shall** [10] be mapped into C_PDU ID NUMBER (field 2).
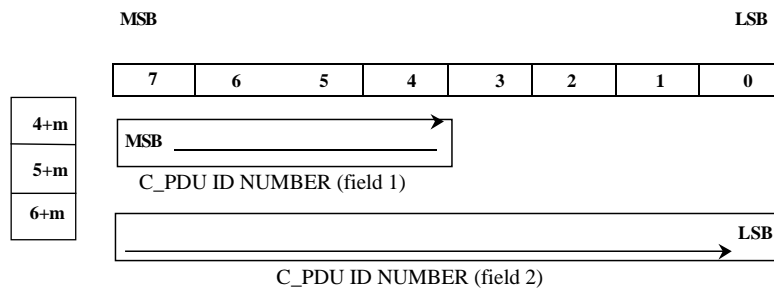

Figure C-26. C_PDU ID NUMBER mapping convention in D_PDU header

The DELIVER IN ORDER flag **shall not** be set.   The bit field is retained for compatibility with Edition 3. Operational experience has shown that in order delivery with the NON-ARQ service is problematic.

If the DELIVER IN ORDER flag is cleared (0) on the D_PDUs composing a C_PDU, the C_PDU **shall** be delivered to the network layer when one of the following condition is met.

1) C_PDU is complete and error free; or
2) The C_PDU RECEPTION WINDOW expires. This can be helpful to partial C_PDU delivery.

The GROUP ADDRESS flag **shall** [13] indicate that the destination address should be interpreted as a group address rather than an individual address, as follows:

- The destination address **shall** [14] be interpreted as a group address when the GROUP ADDRESS flag is set (1).
- However when the GROUP ADDRESS flag is cleared (0) the destination address **shall** [15] be interpreted as an individual node address.

Note: If a specific PDU is intended for only one node, the node's normal address may also be used with the NON-ARQ DATA D_PDU. Group addresses are intended to allow PDUs to be addressed to specific groups of nodes when using NON-ARQ DATA D_PDUs. The use of a bit to designate a "group address" allows the same number (~268 million!) and structure of group addresses to be designated as for normal addresses, rather than requiring some portion of the total address space for group addresses. The management of group addresses is outside of the scope of this STANAG.

The SIZE OF SEGMENTED C_PDU field **shall** [16] specify the number of bytes contained in the SEGMENTED C_PDU file in accordance with the requirements of Section C.4.2.10.

The C_PDU SIZE field **shall** [17] indicate the size in bytes of the C_PDU of which the C_PDU segment encapsulated in this D_PDU is a part.

The value of the C_PDU SIZE field **shall** [18] be encoded in a 16 bit field, with the bits mapped as specified by

Figure C-27. The number **shall** [19] be mapped into the D_PDU by placing the MSB of the field into the MSB of the first byte in the D_PDU as can be seen in Figure C-28.

Figure C-27. C_PDU SIZE Field

Figure C-28. C_PDU SIZE mapping convention in D_PDU header

The C_PDU SEGMENT OFFSET field **shall** [20] indicates the location of the first byte of the SEGMENTED C_PDU with respect to the start of the C_PDU. For the purposes of this field, the bytes of the C_PDU **shall** [21] be numbered consecutively starting with 0.

The C_PDU SEGMENT OFFSET field is a 16 bit field, the bits **shall** [22] be mapped as specified by Figure C-29. The number **shall** [23] be mapped into the D_PDU by placing the MSB of the field into the MSB of the first byte in the D_PDU as specified in Figure C-30.

Figure C-29. C_PDU SEGMENT OFFSET Field

Figure C-30. C_PDU SEGMENT OFFSET mapping convention in D_PDU header

The C_PDU RECEPTION WINDOW field **shall** [24] indicate the maximum remaining time in units of half (1/2) seconds relative to the start of the D_PDU during which portions of the associated C_PDU may be received.

Setting this value needs some care, as if set too large it can delay delivery of partial C_PDUs. If set to low, it will miss valid C_PDU segments.

As in the case of the EOT field, the C_PDU RECEPTION WINDOW **shall** [25] be updated just prior to transmitting each D_PDU. The receiving node can use this information to determine when to release a partially received C_PDU. (The transmitter is not allowed to transmit D_PDUs that are a portion of a C_PDU when the C_PDU RECEPTION WINDOW has expired).

The value of the C_PDU RECEPTION WINDOW field **shall** [26] be encoded in a 16 bit field with the bits be mapped as specified by Figure C-31. The value **shall** [27] be mapped into the D_PDU by placing the MSB of the field into the MSB of the first byte in the D_PDU as specified in Figure C-32.


Figure C-31. C_PDU RECEPTION WINDOW Field

Figure C-32. C_PDU RECEPTION WINDOW mapping convention in D_PDU header

C.4.12. EXPEDITED-NON-ARQ-DATA (TYPE 8) D_PDU (Expedited non-ARQ data transfer)
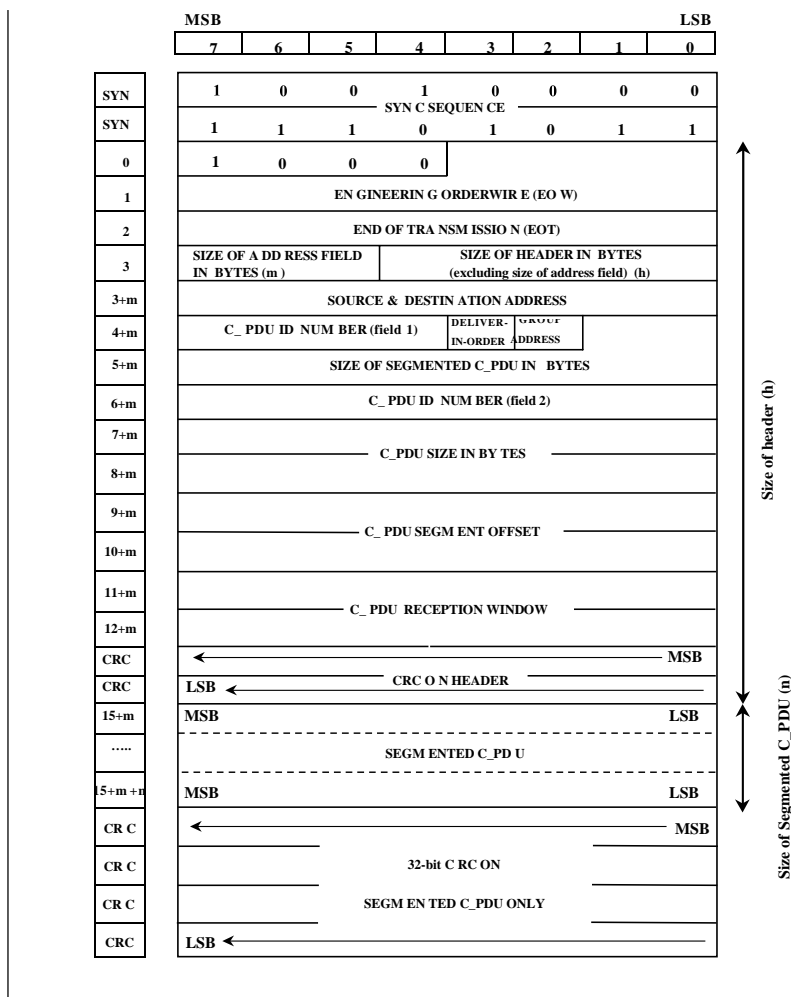
Figure C-33. Frame format for EXPEDITED-NON-ARQ DATA D_PDU Type 8

The frame format for EXPEDITED-NON-ARQ-DATA (TYPE 8) D_PDUs **shall** [1] be identical to the NON-ARQ-DATA D_PDU with the exception that the TYPE field has a value of 8, as specified in Figure C-33.

The C_PDU ID NUMBER space (i.e, the set of ID numbers in the range [0..4095]) for EXPEDITED NON- ARQ DATA (TYPE 8) D_PDUs **shall** [2] be different than the similarly-defined number space for NON- ARQ DATA (TYPE 7) D_PDUs.

C.4.13. EXTENSION (TYPE 13) D_PDU

	MSB 7	6	5	4	3	2	1	LSB 0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m*	MSB	Extended D_PDU Type						LSB
	Optional Header Bytes that may be specified in Extended D_PDU							
CRC CRC	CRC on Header							
h+m	Optional Bytes that may be specified in Extended D_PDU							
CRC CRC CRC CRC	Optional CRC							
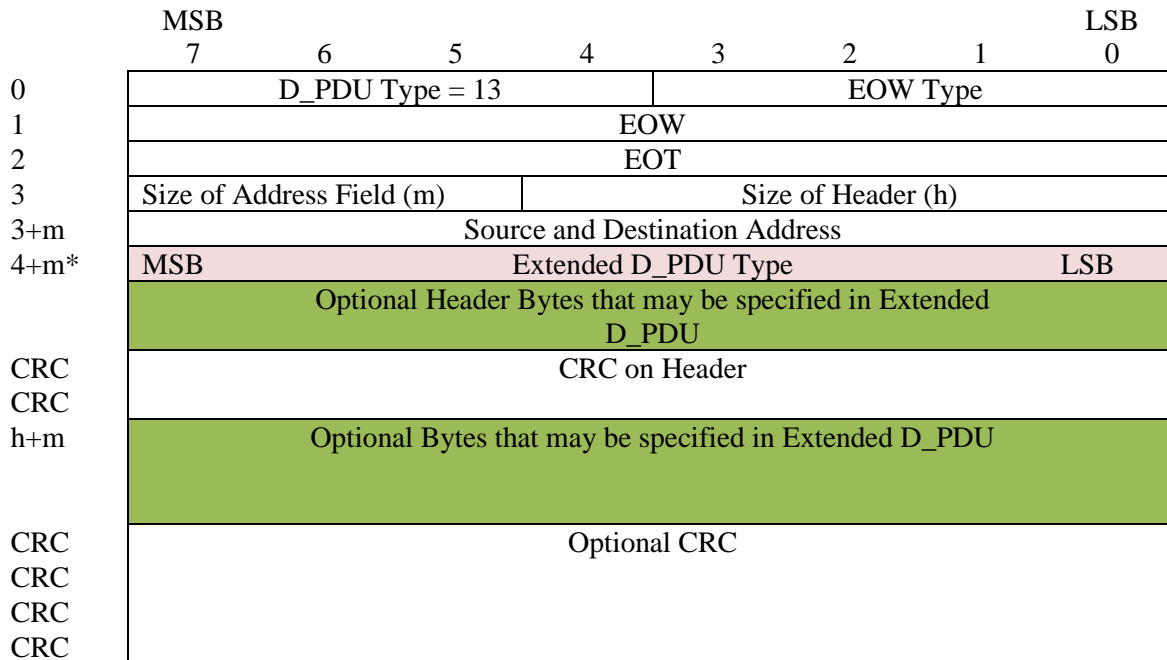
Figure C-34. Frame format for EXTENSION D_PDU Type 13

All of the D_PDU Type values are assigned in this edition of Annex C. The EXTENSION D_PDU (TYPE 13) D_PDU enables future versions of this Annex and other Annexes of STANAG 5066 to define new D_PDU types. A list of assigned Extended D_PDU types is set out in Section C.8.

EXTENSION D_PDU **shall** only be used with Edition 4 (or subsequent) peers.

The EXTENSION D_PDU format is shown in Figure C-34. The Extended D_PDU Type field Extended D_PDU Type allows for up to 256 new D_DPDUs to be defined. Extended D_PDU can have two formats:

1. Header only format, as specified in Figure C-2(a), which ends with two byte CRC on header.
2. The format specified in Figure C-2(b), which had data beyond the header, ending with a four byte checksum on the data beyond the header.

The choice of format **shall** be specified for each Extended D_PDU Type. The format choice **may** vary based on information in the D_PDU header.

C.4.14.   PADDING (TYPE 14) D_PDU

	MSB							LSB
	7	6	5	4	3	2	1	0
0	D_PDU Type = 14				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)				Size of Header (h)			
3+m	Source and Destination Address							
CRC	CRC on Header							
CRC								

Figure C-35. Frame format for PADDING D_PDU Type 14

PADDING D_PDU **shall** only be used with Edition 4 (or subsequent) peers.

The format of the PADDING D_PDU (Type 14) is shown in Figure C-35. This is a minimal D_PDU with no fields specific to the D_PDU type.

The DTS sends out data as stream of D_PDUs.   Modern HF Waveforms send data in blocks, which means that the volume of data sent needs to be an exact number of block sizes.   There is generally some extra space for one of two reasons:

    a.   There is insufficient space to fit in the next D_PDU (often a DATA-ONLY D_PDU which might typically by 2-300 bytes); or
    b.   Only a small amount of data is being sent, and there is space.

This annex allows D_PDUs to be repeated, which is preferable to padding with non-DPDU data.   Benefits of repeating D_PDUs:

1.   If you repeat a D_PDU, risk of data loss is reduced as the original may be corrupted.
2.   An EOW (Engineering Order Wire) single byte message can be sent with each D_PDU to exchange control information. Extra D_PDUs allow for more EOWs to be sent which allows for more information to be sent and reduces the risk of (important) EOW information from being lost.
3.   The EOT (End of Transmission) is a single byte sent with each DPDU to indicate time of transmission remaining.   This allows the receiver to determine the length of a transmission, which will enable it to know when the transmission is complete and when it can start transmission. This is very important for resilient operation. Transmitting more EOTs reduces the risk of all EOT information getting lost.

EOT and EOW repeats are particularly important at slower speeds when only a small number of D_PDUs will be sent in each transmission. It will generally make sense to repeat critical data. ACK-ONLY D_PDUs are particularly useful to repeat and are small. It can also make sense to repeat high priority D_PDUs, giving priority to SAPs with low latency QoS requirements. Note that when a physical link has been established (CAS-1 for ARQ traffic) that ACK-ONLY D_PDSs can be used in both directions which is recommended.

The PADDING D_PDU is provided as a minimal D_PDU to gain the second and third advantages of repeating D_PDUs, when it is not possible to send another D_PDU because all the valid options are too large.   This is

particularly likely to happen for non-ARQ traffic, where it is not possible to send and ACK-ONLY D_PDU. It can also happen where the ACK-ONLY D_PDUs available to transmit are too large.

The PADDING D_PDU is encoded following the generic D_PDU structure as shown in Figure C-2(a) with the D_PDU Type set to 14. There are no D_PDU specific fields for the PADDING D_PDU.

Note that EOWs in a padding D_PDU are directed to the node or nodes indicated by the destination address, which may be a broadcast address.

C.4.15. WARNING (TYPE 15) D_PDU

The WARNING D_PDU is sent in response to unexpected or unrecognised D_PDU type.
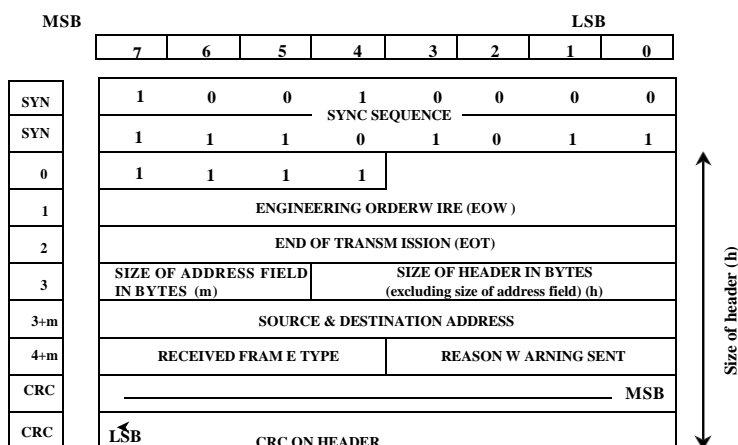
Figure C-36. Frame format for WARNING D_PDU Type 15

A Data Transfer Sublayer **shall** [(1)] a WARNING D_PDU to any remote node from which an unexpected or unknown D_PDU type has been received.

The WARNING D_PDU **shall** [(2)] contain the following fields within its D_PDU Type-Specific part, mapped and encoded in accordance with Figure C-36 and the paragraphs below:

    ☐ RECEIVED FRAME TYPE
    ☐ REASON WARNING SENT

The RECEIVED FRAME TYPE field **shall** [(3)] indicate the frame type that caused the warning to be sent.

The value of the RECEIVED FRAME TYPE field **shall** [4] be encoded in four bits, and located within the D_PDU as specified in Figure C-37 and Figure C-38.

Bit mapping order

MSB   LSB

3	2	1	0

RECEIVED FRAME TYPE

Figure C-37. RECEIVED FRAME TYPE Field

MSB   **LSB**

7	6	5	4	3	2	1	0

4+m	**MSB** ⟶ **LSB**

Figure C-38. RECEIVED FRAME TYPE mapping convention in D_PDU header

The REASON WARNING SENT field **shall** [5] indicate the reason the frame type caused a warning, with values as Specified in Table C-3:

Table C-3: Encoding of WARNING D_PDU Reason Field

Reason	Field Value
Unrecognised D_PDU type Received	0
Connection-related D_PDU Received While Not Currently Connected	1
Invalid D_PDU Received	2
Invalid D_PDU Received for Current State	3
Unspecified/reserved	4-15

The value of the REASON WARNING SENT field **shall** [6] be encoded in four bits, and located within the D_PDU as specified in Figure C-39 and Figure C-40.

Bit mapping order

3	2	1	0

REASON WARNING SENT

Figure C-39. REASON WARNING SENT Field

MSB   **LSB**

7	6	5	4	3	2	1	0

4+m   **MSB**   **LSB**

Figure C-40. REASON WARNING SENT mapping convention in D_PDU header

The transmission of WARNING type D_PDUs **shall** [7] be initiated independently by the Data Transfer Sublayer in response to certain D_PDUs and **shall** [8] not be acknowledged explicitly.

WARNING type D_PDUs may be inserted into an ongoing transmission interval provided that the transmission interval period is not exceeded.

A WARNING D_PDU **shall** [9] be sent in the following conditions:
1. A node receives a D_PDU header addressed to itself with a valid CRC and an unrecognised D_PDU type or unrecognized Extended D_PDU Type (value 0000)
2. A node is not in the IDLE/BROADCAST state and it receives a D_PDU header addressed to itself, from a node with which it is not currently connected. (value 0001)
3. A node is in IDLE/BROADCAST state and it receives a D_PDU header addressed to itself which is other than type 7 or type 8 D_PDU (value 0010)
4. A node receives any D_PDU which is recognized but is not of the allowed type for the state which the receiving node is in (value 0011; this is the general case of the preceding)

A WARNING D_PDU **shall** [10] not be sent in response to receipt of a WARNING D_PDU.

C.5.   C_PDU Segmentation and Re-assembly Processes

The process of C_PDU segmentation and re-assembly **shall** [(1)] be as defined in the subsections that follow for ARQ and non-ARQ delivery services provided to regular and expedited C_PDUs.

C.5.1.   ARQ Mode Segmentation and Re-assembly
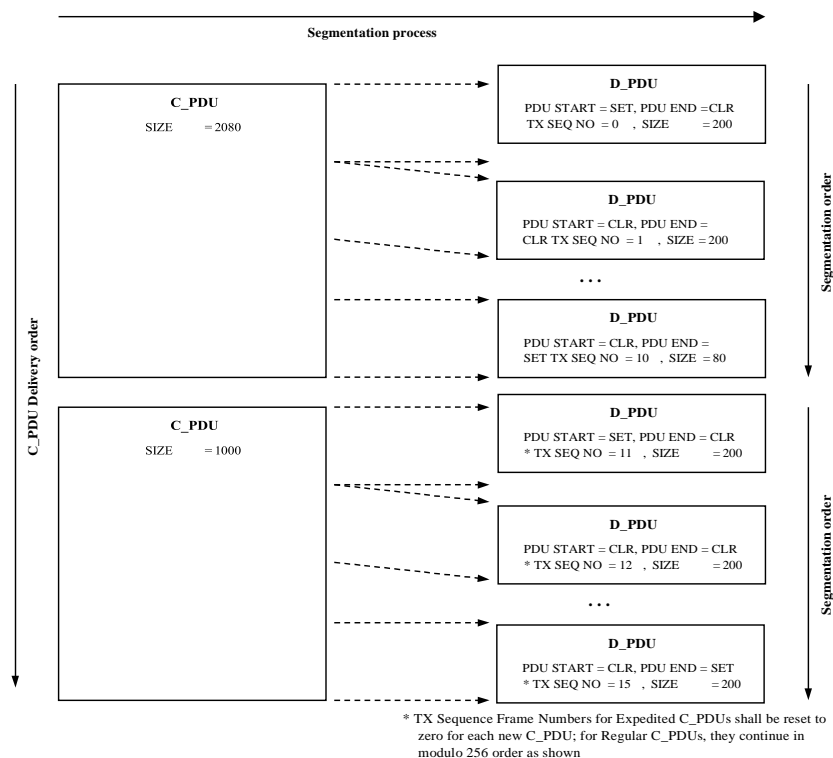
Figure C-41. C_PDU Segmentation for ARQ Delivery Services (Regular and Expedited Service)

Segmentation of a C_PDU into segments small enough to fit within a D_PDU for ARQ-delivery (i.e, a DATA, DATA-ACK, or EXPEDITED-DATA D_PDU) **shall** [(1)] be performed in accordance with the example shown in Figure C-41 and as follows:

   a.   The Maximum C_PDU-Segment Size within a D_PDU for ARQ-Delivery services **shall** be a configurable parameter no greater than 1023 bytes in any implementation compliant with this STANAG. An implementation may configure the Maximum C_PDU-Segment Size to match the interleaver size for optimum channel efficiency or other reasons.

   b.   An entire C_PDU that is smaller than the Maximum C_PDU-Segment Size **shall** [(3)] be placed in the C_PDU segment of a single D_PDU.

   c.   A DATA or DATA-ACK, or EXPEDITED D_PDU that contains an entire C_DPU **shall** [(4)] be marked with the both C_PDU START field and the C_PDU END field set

equal to the value "1". [Note: An 'only' C_PDU segment is both the "first" and "last" segment of a sequence of one.]

d. The Data Transfer Sublayer **shall** [5] divide C_PDUs larger than the Maximum C_PDU-Segment Size into segments that are no larger than the Maximum C_PDU Segment Size.

e. Only the last segment or the only segment taken from a C_PDU may be smaller than the Maximum C_PDU-Segment size. A C_PDU smaller than the Maximum C_PDU-Segment size **shall** [6] be placed only in the D_PDU that contains the last segment of the C_PDU, i.e., only in a D_PDU for which the C_PDU END field is set equal to one.

f. The bytes within a C_PDU segment **shall** [7] be taken from the source as a contiguous sequence of bytes in the same order in which they occurred in the source C_PDU.

g. D_PDUs containing C_PDU segments taken in sequence from the source C_PDU **shall** [8] have sequential Frame Sequence number fields, modulo 256. [Note: With the first C_PDU segment placed in a D_PDU with Frame Sequence = P, the second would have Frame Sequence = P+1, the third P+2, and so on, with Frame-Sequence operations performed modulo-256.]
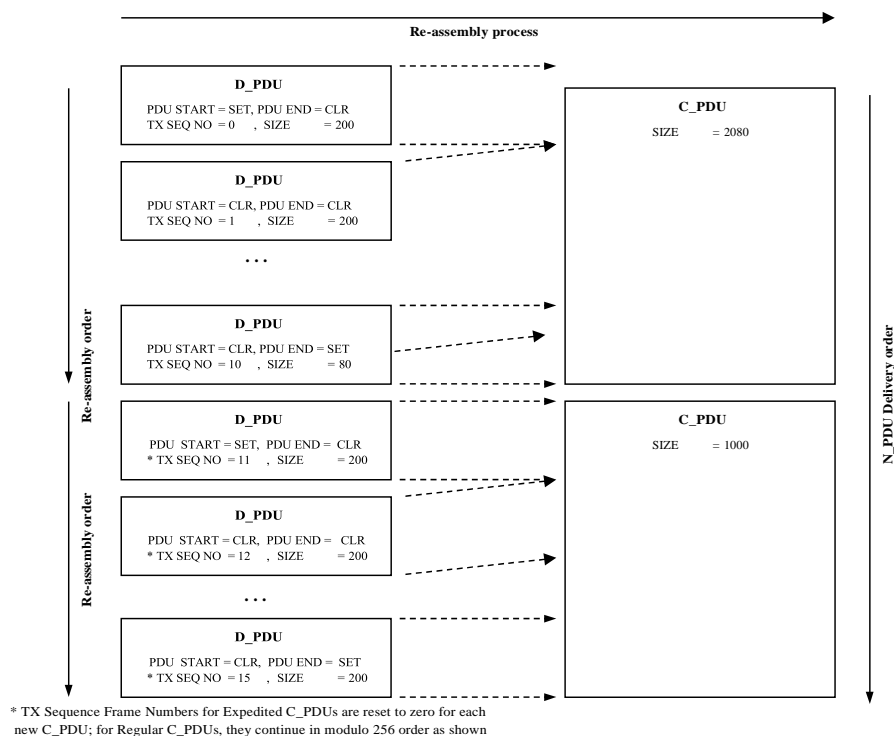

Figure C-42. C_PDU Re-assembly for ARQ Delivery Services (Regular and Expedited Service)

Re-assembly of a C_PDU from its segments **shall** [9] be performed in accordance with the example shown in Figure C-42 and as follows (unless noted otherwise, C_PDU segments that are reassembled are expected to have passed the CRC error-check and have no detectable errors):

(a)   The re-assembly process for C_PDUs receiving ARQ-service **shall** [9] use the Frame-Sequence-Number field, C_PDU START flag, and C_PDU END flag to determine when all segments of a C_PDU have been received.

(b)   A C_PDU segment taken from a D_PDU whose C_PDU START and C_PDU END flags are both set to the value "one" (1) **shall** [10] be taken as a single C_PDU and processed as follows:

   1)   If the D_PDU is for a regular (unexpedited) data type, the C_PDU **shall** [11] be delivered to the Channel Access Sublayer using a D_UNIDATA_INDICATION primitive;

   2)   If the D_PDU is for an unexpedited data type, the C_PDU **shall** [12] be delivered to the Channel Access Sublayer using a D_EXPEDITED_UNIDATA_INDICATION primitive;

(c)   A segment from a C_PDU larger than the Maximum C_PDU Segment Size **shall** [13] be combined in modulo 256 order with other segments whose D_PDU Frame Sequence Numbers lie in the range defined by the Frame Sequence Numbers of the C_DPU START and C_PDU END segments;

(d)   A completely reassembled C_PDU **shall** [14] be delivered to the Channel Access Sublayer using the appropriate D_Primitive.

C.5.1.1. Notes on ARQ Selection of Maximum C_PDU Segment Size

The best choice of Maximum C_PDU Segment Size is related to the choice of transmission speed and interleaver. Often, particularly at lower HF speeds, it will be desirable to optimize throughput.   When optimizing throughput, high Frame Error Rates (FER) which may be around 50% FER will lead to best throughput.   At higher HF and WBHF speeds, it will sometimes be preferable to optimize for low latency, leading to selection of a more conservative transmission speed.

When optimizing for throughput, the choice of Maximum C_PDU Segment Size is critical. If too small a size is chosen, the overhead of D_PDU headers is too large. If too large a size is chosen, the overhead of retransmission due to frame errors is too large.

STANAG 5066 Edition 3 Annex H (Implementation Notes and Guidance) notes various research on measurements at 600 – 1200 bps, looking to optimize throughput:

*   SHAPE Technical Center (STC) - "Laboratory and Field Tests of the High Frequency OSI Data Link Protocol", STC TN-506, August 1993.
*   Implementation experience from Rohde & Schwarz and DRA
*   NATO NC3A TM-937

A key observation is that a Maximum C_PDU Segment Size of 200 bytes is a "good compromise".

It is also noted "If there is a desire to vary frame size in some way, STANAG 5066 supports the use of variable

frame sizes. While the data available to date suggest that the benefit may be marginal, it would be possible, for example, to associate a certain frame size with each data rate."

More recent observations on this choice have been made by Isode and reported on in various papers including " Optimizing applications and data links for HF radio intermediate term variations: Can you ride the wave?", Steve Kille, Nordic HF 2016. This suggests:

- At mid-range narrowband HF speeds, a Maximum C_PDU Segment Size of 200-300 bytes is a good choice. This is in line with the earlier observations.
- When optimizing for throughput, long transmissions give the best performance.
- Intermediate Term Variation (10-120 secs) gives key impact on choice of Maximum C_PDU Segment Size.
- At 75 or 150 bps, a Maximum C_PDU Segment Size of 100 bytes is a good choice.
- At speeds increase into WBHF range increasing Maximum C_PDU Segment Size to larger values up to the maximum of 1023 bytes is optimal.

Collectively, these results suggest that selection of Maximum C_PDU Segment Size based on anticipated data rate is desirable.

C.5.2.   NON-ARQ Mode Segmentation and Re-assembly
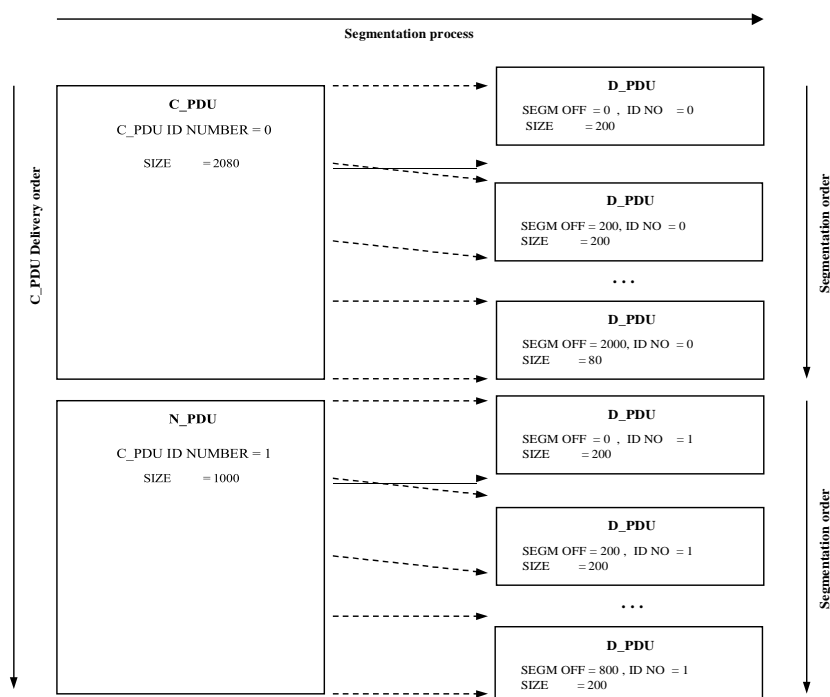
Figure C-43. C_PDU Segmentation for Non-ARQ Delivery Services (Regular and Expedited Service)

Segmentation of a C_PDU into segments small enough to fit within a D_PDU for non-ARQ-delivery (i.e, a Non-ARQ DATA, or EXPEDITED-Non-ARQ-DATA D_PDU) **shall** [(1)] be performed in accordance with the example shown in Figure C-43, and as follows:

a. The Maximum C_PDU-Segment Size within a D_PDU for non-ARQ-Delivery services shall (2) be a configurable parameter no greater than 1023 bytes in any implementation compliant with this STANAG. An implementation may configure the Maximum C_PDU-Segment Size to match the interleaver size for optimum channel efficiency or other reasons;

b. An entire C_PDU for non-ARQ delivery that is smaller than the Maximum C_PDU-Segment Size **shall** [(3)] be placed in the C_PDU segment of a single D_PDU;

c. A unique C_PDU ID number **shall** [(4)] be assigned to the non-ARQ C_PDU in accordance with the requirements of Section C.4.11;

d. all D_PDUs containing segments from the same C_PDU **shall** [(5)] have the same C_PDU ID number;

e. The Segment Offset field of the D_PDU containing the first segment from a C_PDU **shall** [(6)] be equal to zero;

f. The Segment Offset field of the D_PDU containing any subsequent segment from a C_PDU **shall** [(7)] be set equal to the number of bytes from the original C_PDU that precede the first byte of the segment.
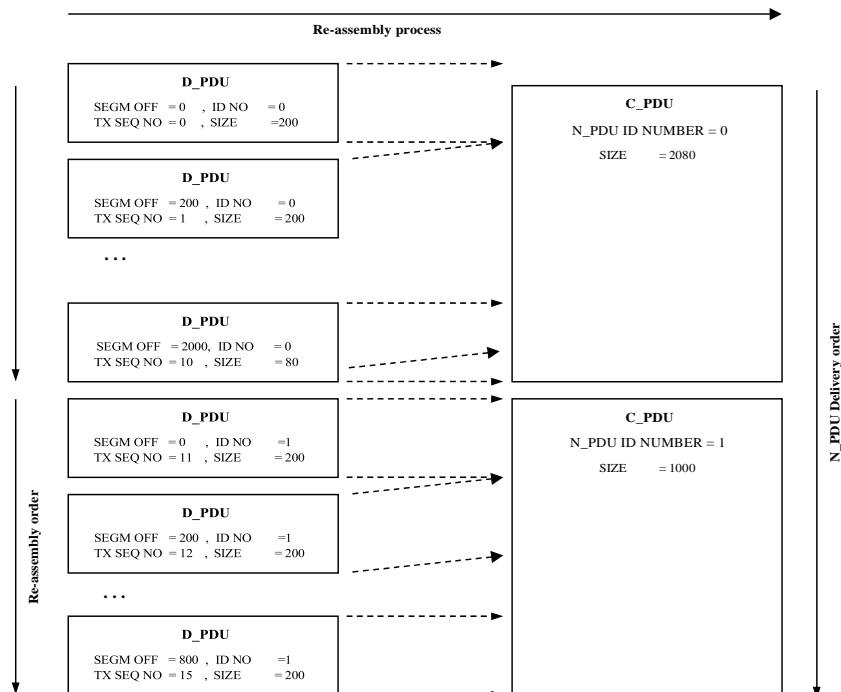
Figure C-44. C_PDU Re-assembly for Non-ARQ Delivery Services (Regular and Expedited Service)

For Non-ARQ services, re-assembly of a C_PDU from its segments **shall** [8] be performed in accordance with the example shown in Figure C-44 and as follows (unless noted otherwise, C_PDU segments that are reassembled are expected to have passed the CRC error-check and have no detectable errors):

a) The re-assembly process for Non-ARQ C_PDUs **shall** [9] use the C_PDU ID Number, Segment-Offset field, C_PDU-Segment-Size field, and C_PDU-Size field to determine when all segments of a C_PDU have been received.

b) If the Error-Free Delivery Mode has been specified, a reassembled C_PDU **shall**[10] be delivered if and only if all segments of the C_PDU have been received without errors;

c) If the Deliver-w/-Errors Mode has been specified, the re-assembly process **shall** [11] proceed as follows:
   1) C_PDU segments received without detected errors **shall** [12] be collected as received in their D_PDUs and placed in order within the reassembled C_PDU;
   2) C_PDU segments received with detected errors **shall** [13] be placed within the reassembled C_PDU just as they are received in their D_PDUs (i.e, with errors), with the size in bytes and the position of the first byte of the segment noted in the D_Primitive used to deliver the C_PDU to the Channel Access Sublayer;
   3) At the end of a specified and configurable timeout-interval the size in bytes and the position of the first byte of any C_PDU segments that have been lost or still not received **shall** [14] be noted in the D_Primitive that delivers the C_PDU to the

ANNEX C
STANAG 5066
Edition 4

**NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW**

Channel Access Sublayer

d) C_PDUs shall be delivered in the order they arrive. Deliver-In-Order Mode shall not be supported for this edition of STANAG 5066.

e) C_PDUs **shall** [16] be delivered to the Channel Access Sublayer as soon as all segments have been received (in Error-Free Mode) or received and accounted for (in Deliver-with-Errors Mode).

f) Delivery of the reassembled D_PDU **shall** [17] be performed with the D_Primitive appropriate for the type of data (i.e., regular or expedited) received.

C.5.2.1. Non-ARQ D_PDU Ordering and Priority

For ARQ service, C_PDU fragments **shall** be transmitted in order with sequential frame sequence numbers. There is no equivalent requirement for non-ARQ, a C_PDU fragments use an independent identifier for correlation.

When transmitting a sequence of non-ARQ C_PDU fragments, it **may** be necessary or desirable to insert other D_PDUs in between them, in particular higher priority D_PDUs arriving after transmission began.

C.5.2.2. Notes on NON-ARQ Selection of Maximum C_PDU Segment Size

When using NON-ARQ it will generally be desirable to use more conservative transmission rates than for ARQ. In some cases (e.g., broadcast), errors will not be corrected and in other use cases correction will be controlled by application timers which will take longer. Because of this, it is important that data loss is low, and this means low transmission rate.

A related consideration is that Non-ARQ allows the sender to request multiple transmissions. If data is sent only once, there no benefit arising from reducing C_PDU Segment Size and a large value (1023 bytes) can be used to reduce overhead. When there are repeats, a smaller Maximum C_PDU Segment Size may be beneficial to minimizing data loss.

For applications using Non-ARQ with errors, smaller Maximum C_PDU Segment Size is potentially beneficial to give the applications error-free fragments. There is little operational experience with such applications.

Note the considerations for Non-ARQ applications are significantly different and can vary between the different classes of Non-ARQ applications noted below. A clear Quality of Service model is desirable to support this.

C.5.2.2.1. **Considerations for Broadcast Applications**

Broadcast applications will either have very slow error handling or none at all for EMCON reception. It is therefore important to choose transmission speed, interleaver, Non-ARQ repeat count and Maximum C_PDU Segment Size with care.

For EMCON applications, it may be appropriate to consider use of non-ARQ with errors.

C.5.2.2.2. **Considerations for Multicast Applications**

For multicast applications, such as ACP 142, there will be application level retransmission of errors and/or some tolerance of loss due to application-level Forward Error Correction.

It is anticipated that for such applications that transmissions rate will be chosen so that base frame error rate is low and that repeating the Non-ARQ transmissions will be sub-optimal. Therefore Maximum C_PDU Segment size is best set to 1023 bytes. If may be desirable to have smaller D_PDUs, which is best achieved by the multicast application choosing to use a smaller APDU size, which will constrain D_PDU size.

C.5.3. Configuration of Maximum C_PDU Segment Size

This annex clearly requires that Maximum C_PDU Segment Size is configurable. It could be interpreted that this is a single fixed value, although notes in Edition 3 Annex H suggest otherwise. This section clarifies interpretation for use with both Edition 3 and Edition 4 peers.

For a given transmission, the Maximum C_PDU Segment Size used to generate new D_PDUs **shall not** vary. The value used **may** vary between transmissions, noting considerations on choice in Section C.5.1.1 and Section C.5.2.1.

At 75bps a D_PDU **may** take 110 seconds to transmit. Any constraints on maximum transmission time **shall not** lead to preventing transmission of a single D_PDU that needs a longer transmission time.

C.5.4. Relaxation of Constraint on C_PDU Segment Size

There is a constrainta on handling segmentation and re-assembly that is relaxed in Edition 4. An implementation conforming to this specification **shall** handle this constraints on reception. An implementation **may** make use of this relaxation when sending to an Edition 4 (or subsequent) peer, but **shall not** use this constraint relaxation when sending to an Edition 3 peer.

When sending any D_PDU, the size **shall not** be greater than the Maximum C_PDU Segment Size but **may** be less than the Maximum C_PDU Segment Size. This relaxation of the Edition 3 constraints gives three benefits:

1. It allows space at the end of a transmission to be filled with a DATA-ONLY D_PDU, this giving better link utilization.
2. It allows a C_PDU to be split into segments of equal size, which will give better performance.
3. It allows D_PDUs to be aligned to waveform block boundaries. When there is fading, this often leads to blocks being badly corrupted. By ensuring that all D_PDUs are in a single modem block, performance is improved.

C.6.   EOW and Management Message Types

The set of EOW Messages listed in this section **may** be placed in the 12-bit EOW section of any D_PDU. These messages are transmitted as information only, with no acknowledgement.

There is also a framework for acknowledged EOWs. This framework was used by EOWs in earlier version of this standard.   Acknowledged EOWs are sent in the MANAGEMENT D_PDU.

The types of EOW used in STANAG 5066 are listed in Table C-4:

Table C-4. EOW Message Types

EOW Type	EOW Name	Description and Reference	Edition
0	EMPTY	Null EOW specified in Section C.6.1.	All
1	ED3-BASIC-RATE	Requesting basic rate, specified in Section C.6.2.	Ed3
2	RESERVED		-
3	UNRECOGNIZED-TYPE-ERROR	Error specified in Section C.6.3	All
4	ED3-CAPABILITY	Capability described in Section C.6.4	Ed3
5	RESERVED		-
6	RESERVED		-
7	RESERVED		-
8	MAX-SPEED	Speed recommended for maximum throughput specified in Section C.6.5.	Ed4
9	LOW-SPEED	Speed recommended for low latency data specified in Section C.6.7	Ed4
10	MAX-SEGMENT	Recommended maximum C_PDU Segment Size specified in Section C.6.7.	Ed4
11	SENDER-APPROACH	Approach used by sender to guide receiver specified in Section C.6.8.	Ed4
12	SPEED-USED	Communicates speed used by sender specified in Section C.6.9.	Ed4
13	TOKEN	Sends Token in WTRP specified in Annex L.	Ed4
14	TOKEN-CRITICAL	Sends Token in WTRP indicating critical Transmit Order List change specified in Annex L	Ed4
15	CAPABILITY	Capability described in Section C.6.10.	Ed4

The first column indicates the EOW number.   The second column defines an EOW Name, that is used to refer to this EOW.   The third column gives a brief description of the EOW and a reference to where the EOW is specified.   The fourth column indicates the use based on peer support of STANAG 5066 Editions:

- "All" **may** be used with any peer.
- "Ed3" **may** be used with any peer known to support Edition 3 (or earlier). It **shall not** be used with any peer known to support Edition 4 (or later) or where peer support is unknown.
- Ed4" **may** be used with any peer known to support Edition 4 (or later). It **shall not** be used with any peer known to support Edition 3 (or earlier) or where peer support is unknown.

It is anticipated that support for the Ed3 EOWs will be dropped in a future edition of STANAG 5066.

Three EOWs are marked as RESERVED. These are EOWs used in previous editions of STANAG 5066. EOW 2 was used for a data rate change procedure for non-auto-baud waveforms that is not supported in this edition. EOW 7 was specified as an alternate rate change extended EOW, that is not widely used. EOWs 5 and 6 were used in Annex I, which has been removed from this edition. It is anticipated that these EOWs may be re-assigned in a future edition, if requirements for more EOWs arise.

The format of the EOW message types **shall** [3] be as shown in Figure C-46.

MSB			LSB	MSB							LSB
11	10	9	8	7	6	5	4	3	2	1	0
TYPE				**contents**							

Figure C-45. Format of EOW Messages

The TYPE field of the EOW message **shall** [4] be filled with the hexadecimal value of the appropriate message type (units only), with the LSB of the TYPE value placed in the LSB of the TYPE field.

The Contents field **shall** [5] be EOW type-specific, in accordance with the subsections below.

C.6.1.   EMPTY (TYPE 0) EOW Message

The EMPTY EOW (Type 0) has the contents set to all zero. This provides a default EOW to use when there are no other EOWs to communicate.

For Edition 4 peers, it is recommended to use the CAPABILITY EOW rather than EMPTY.

C.6.2. ED3-BASIC-RATE (TYPE 1) EOW Message

MSB			LSB	MSB			LSB	MSB	LSB	MSB	LSB
11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	1	Data Rate				Interleaving		Other parameters	
	TYPE										

Figure C-46. Format of ED3-BASIC-RATE EOW Message.

The ED3-BASIC-RATE (TYPE 1) EOW Message **shall** [1] be used in conjunction with the Data Rate Selection Procedure, as specified in Section C.7.6.

The ED3-BASIC-RATE (TYPE 1) EOW Message **shall** be used by a receiving node as an advisory message indicating the modem parameters to which the link may be set to provide optimum performance.

The ED3-BASIC-RATE (TYPE 1) EOW Message **shall** [2] be formatted and encoded as specified in Figure C-46 and the paragraphs that follow, and includes the following type-specific subfields:
- Data Rate
- Interleaving
- Other Parameters

The Data Rate parameter **shall** [3] be the rate at which the node originating the message recommends that the peer transmit data, in accordance with the encoding defined in the following table:

Table C-5.   Data Rate Parameter ED3-BASIC-RATE

MSB - LSB	Interpretation
0 0 0 0	75 bps
0 0 0 1	150 bps
0 0 1 0	300 bps
0 0 1 1	600 bps
0 1 0 0	1200 bps
0 1 0 1	2400 bps
0 1 1 0	3200 bps
0 1 1 1	3600 bps
1 0 0 0	4800 bps
1 0 0 1	6400 bps
1 0 1 0	8000 bps
1 0 1 1	9600 bps
1 1 0 0	14400 bps
1 1 0 1	16000 bps
1 1 1 0	19200 bps
1 1 1 1	reserved

The Interleaver Parameter field **shall** [4] specify the interleaver requested for use by the node producing the message (for that link, if multiple links/modems are in use) with respect to transmit and receive operation, in

accordance with the following table:

Table C-6. Interleaver Parameter: ED3-BASIC-RATE

MSB - LSB	Interpretation
0 0	no interleaving
0 1	short interleaving
1 0	long interleaving
1 1	reserved

The Other Parameters field **shall** [5] specify the capabilities of the modem in use by the node producing the message (for that link, if multiple links/modems are in use) with respect to transmit and receive data rates, and whether the message is an advisory message or request message, in accordance with the following table:

Table C-7. Contents for ED3-BASIC-RATE (Other Parameters)

MSB - LSB	Interpretation
0 0	DRC Request: master has independent data rate (change applies to Tx data rate only)
0 1	DRC Request: Tx and Rx data rate at master must be equal (change will apply to both Tx and Rx data rates)
1 0	DRC Advisory: Advising node has independent data rate for Tx and Rx (change applies to Rx data rate only)
1 1	DRC Advisory: Tx and Rx data rate at advising node must be equal (change will apply to both Tx and Rx data rates)

The encodings of Other Parameters is included for reference only. This field **shall** be set to "1 0" when transmitting this EOW and ignored when receiving it.

C.6.3. UNRECOGNIZED-TYPE-ERROR (TYPE 3) EOW Message

MSB 7	6	5	4	3	2	1	LSB 0
Reserved for Future Use				Unrecognized EOW Type			

Figure C-47. Format of UNRECOGNIZED-TYPE-ERROR EOW Message

The UNRECOGNIZED-TYPE ERROR (Type 3) EOW Message **shall** [1] be used to declare an error related to receipt of an EOW message.

A node **should** send an UNRECOGNIZED-TYPE ERROR (Type 3) EOW Message to the originator of an unrecognised EOW message whenever the node receives an EOW that it does not recognize.

The UNRECOGNIZED-TYPE ERROR (Type 3) EOW Message **shall** [2] be encoded as shown in Figure C-47 and include the following field:
* Unrecognized EOW Type

The type of the unrecognised EOW message or the message that triggered the error **shall** [3] be placed in the Unrecognized EOW Type.

The bits reserved for future use **shall** be set to the value zero (0).

C.6.4. ED3-CAPABILITY (TYPE 4) EOW Message

MSB			LSB	MSB							LSB
11	10	9	8	7	6	5	4	3	2	1	0
TYPE				contents							

Figure C-48. Format of ED3-CAPABILITY EOW Message

The ED3-CAPABILITY Type 4 allows a node to inform another node of its capabilities related to modem parameters, waveform, and control.

This capability EOW is retained to enable capability exchange with an Edition 3 system. Capability exchange with an Edition 4 (or later) shall use the CAPABILITY (Type 15) EOW.

The ED3-CAPABILITY Type 4 **shall** [1] be encoded as shown in Figure C-48 and contains a single field, Contents.

The Contents field **shall** [2] be encoded as the bit-mapped specification of capabilities defined in the following table:

Table C-8. Contents of Message field for ED3-CAPABILITY

bit	meaning
7 (MSB)	Adaptive modem parameters (DRC) capable [note 1] (0 = no, 1 = yes)
6	STANAG 4529 available [note 2] (0 = no, 1 = yes)
5	MIL-STD-188-110 75-2400 bps available [note 2] (0 = no, 1 = yes)
4	extended data rate capable [note 3] (0 = no, 1 = yes)
3	full duplex supported [note 4] (0 = no, 1 = yes)
2	split frequency supported [note 4] (0 = no, 1 = yes)
1	non-ARCS ALE capable [note 5] (0 = no, 1 = yes)
0 (LSB)	ARCS capable [note 6] (0 = no, 1 = yes)

Notes
1. If a node is DRC capable, it must implement at minimum 75 bps through 2400 bps (1200 bps for STANAG 4529) with short and long interleaving in accordance with the relevant document.
2. All nodes shall have, at minimum, the STANAG 4285 waveform available.
3. Annex G describes waveforms for data rates above 2400 bps.
4. A full duplex node must have split frequency operation available.
5. non-ARCS ALE capability may imply MIL-STD-188-141 Appendix A or proprietary capabilities and any ambiguity must be resolved outside of STANAG 5066.
6. ARCS –capable systems are those equipped with NATO's Automatic Radio Control System for HF Links, STANAG 4538 (FLSU/RLSU) or MIL-STD-188-141B (RLSU); any ambiguity must be resolved outside of STANAG 5066.

C.6.5.   MAX-SPEED (TYPE 8) EOW Message

The MAX-SPEED (TYPE 8) EOW is sent by a node receiving data from a peer node to communicate to that node its recommendation for transmission speed to be used to obtain maximum throughput.   It is likely that transmission at this speed will lead to a significant fraction of transmitted D_PDUs being corrupted and subsequently retransmitted.

The MAX-SPEED (TYPE 8) EOW is encoded as specified in Section C.6.5.1, which encodes a transmission speed and interleaver.   The transmission speed is the recommended speed to achieve maximum throughput.

The Interleaver field specifies the minimum length Interleaver recommended for use with this transmission speed. Longer interleavers give better performance for data, and so it will often be appropriate for the sender to select a longer interleaver.

C.6.5.1. Transmission Speed and Interleaver Encoding

Several EOWs need a single byte representation of transmission speed and interleaver, so this is specified in an independent section.   The encoding is shown in Figure C-49.

MSB 7	6	5	4	3	2	1	LSB 0
Interleaver			Speed				

Figure C-49. Transmission Speed and Interleaver Encoding.

Figure C-49   defines a one byte encoding of transmission speed and interleaver.   Transmission speed is represented as a five bit Speed field and Interleaver as three bit Interleaver field.   The values of these fields are specified in Figure C-50 and Figure C-51.

Value	Speed (bps)
0001-1101	WBHF. Waveform is the STANAG 5069 Waveform number (1-13)
10001	75
10010	150
10011	300
10100	600
10101	1200
10110	2400
10111	3200
11000	4800
11001	6400
11010	8000
11011	9600
11100	12800
Other	Reserved

Figure C-50. Speed Encoding of Transmission Speed/Interleaver

The Speed field specified transmission speed as shown in Figure C-50.   This uses two ranges:

1. Values 00001-1101 (decimal 1-13) are used for Wideband HF (WBHF) waveforms as specified in STANAG 5069.   STANAG 5069 defines 13 waveforms, with the waveform directly identified by the Speed value.   Bandwidth (3kHz – 48 kHz) will be fixed for a sequence of transmissions, usually negotiated by 4G ALE.   It is expected that this field will be generated and interpreted in context of the current bandwidth,
2. Values 10001-11100 are used for narrowband HF, and represent each of the standard speeds.

Value	Interleaver
000	No recommendation
001	Ultra Short
010	Very Short
011	Short
100	Medium
101	Long
110	Very Long
111	Reserved

Figure C-51. Interleaver Encoding of Transmission Speed/Interleaver

The Interleaver field is encoded as shown in Figure C-51, with the specified bits representing the chosen interleaver.

C.6.6. LOW-SPEED (TYPE 9) EOW Message

The LOW-SPEED (TYPE 9) EOW is sent by a node receiving data from a peer node to communicate to that node its recommendation for maximum transmission speed to be used to used for data where low latency is desired. This is a speed where it is estimated by the receiver that only a very small fraction of D_PDUs will be corrupted on transmission.

The LOW-SPEED (TYPE 9) EOW is encoded as specified in Section C.6.5.1, which encodes a transmission speed and interleaver.

The recommended maximum transmission speed for low latency data is encoded directly. An interleaver is also encoded, which is the minimum length of interleaver recommended for low latency data.

C.6.7. MAX-SEGMENT (TYPE 10) EOW Message

The MAX-SEGMENT (TYPE 10) EOW is sent by a node receiving data from a peer node to communicate to that node its recommendation for the maximum C_PDU segment size. This is particularly useful for fixed speed transmission, where reducing C_PDU segment size is the only option to deal with a poor quality link. It can also provide supplementary information for use in conjunction with MAX-SPEED EOW.

The MAX-SEGMENT (TYPE 10) EOW is encoded as an integer (range 0-255). The recommended maximum CPDU segment size is determined by multiplying this integer by four (4).

C.6.8.   SENDER-APPROACH (TYPE 11) EOW Message

The SENDER-APPROACH (TYPE 11) EOW is sent by a node transmitting data ta a peer node to communicate to that node information on its approach to sending data.   This information can enable the receiving node to communicate back information that is most helpful to the sending node. The choices are discussed in Section C.7.6. This section defines the encoding.

MSB 7	6	5	4	3	2	1	LSB 0
Reserved for Future Use				EOW 12 Needed	Strategy		Fixed

Figure C-52. Format of SENDER-APPROACH EOW Message

The encoding of the SENDER-APPROACH (TYPE 11) EOW Message is shown Figure C-52.

- Fixed is bit 0. If Fixed is set to 1, transmission **shall** be at fixed speed.   If Fixed is set to 0, transmission speed **may** be varied.

- Strategy (Bits 1 and 2) have the following meanings:
    o   00: Bulk data being transferred (to be optimized for throughput).
    o   01: Low latency data being transferred.
    o   10: A mix of bulk and low latency data being transferred.
    o   11: Bulk data with low latency requirements (e.g., Web browsing).

- If EOW 12 Needed (Bit 3) is set, this indicates that the local system cannot determine modem speed and/or interleaver of received transmissions. This **may** be treated by the receiver as a request to send SPEED-USED (TYPE 12) EOW Messages to the node. If EOW 12 Not Needed is not set the receiver should not send any SPEED-USED (TYPE 12) EOW Messages to the node.

- Bits 4-7 are reserved for future use and **shall** be set to 0.

C.6.9.   SPEED-USED (TYPE 12) EOW Message

The SPEED-USED (TYPE 12) EOW is sent by a node transmitting data to a peer node to communicate the transmission speed and interleaver being used in the current transmission.   This is to provide this information when the receiver cannot determine this information locally.

The SPEED-USED (TYPE 12) EOW is encoded as specified in Section C.6.5.1, which encodes a transmission speed and interleaver.

C.6.10.   CAPABILITY (TYPE 15) EOW Message

The CAPABILITY (TYPE 15) EOW is sent by a node to communicate to peers its capabilities.

MSB 7	6	5	4	3	2	1	LSB 0
Reserved for Future Use						Duplex	Ed4

Figure C-53. Format of CAPABILITY EOW Message

The format of CAPABILITY (TYPE 15) EOW is shown in Figure C-53.   The

- The Ed4 bit (bit 0) **shall** be set if the node supports Edition 4 of STANAG 5066 (or subsequent). Use of this information is summarized in Section C.3 "Support for Edition 3 Interoperability".

- The Duplex bit (bit 1) is set if the node can support duplex capabilities with separate transmit and receive channels.

- Bits 2-7 are reserved for future use and **shall** be set to 0.

C.7. Peer-to-peer Communication Protocols

This section discusses the interactions between the Data Transfer Sublayer entities at different nodes in terms of states, state-transition diagrams, and state tables for a hypothetical state machine. This STANAG does not mandate a state machine implementation. The requirement for interoperability is that the system act consistently with the state-transition and actions rules for message exchange and format presented in this STANAG.

C.7.1. Data Transfer Sublayer States and Transitions

The expected and allowed interactions of one node with another are described herein with respect to states of the node's Data Transfer sublayer. Receiving certain PDUs (from the modem) or D_Primitives (from the Channel Access Sublayer) will cause a node, depending on its state, to transmit certain PDUs and/or D_Primitives, and/or change to another state.

The Data Transfer Sublayer interactions with peers **shall** [1] be defined with respect to the states shown in Figure C-54 and as follows:

IDLE(UNCONNECTED)	IDLE(CONNECTED)
DATA(UNCONNECTED)	DATA(CONNECTED)
EXPEDITED-DATA(UNCONNECTED	EXPEDITED-DATA(CONNECTED)
MANAGEMENT(UNCONNECTED)	MANAGEMENT(CONNECTED)
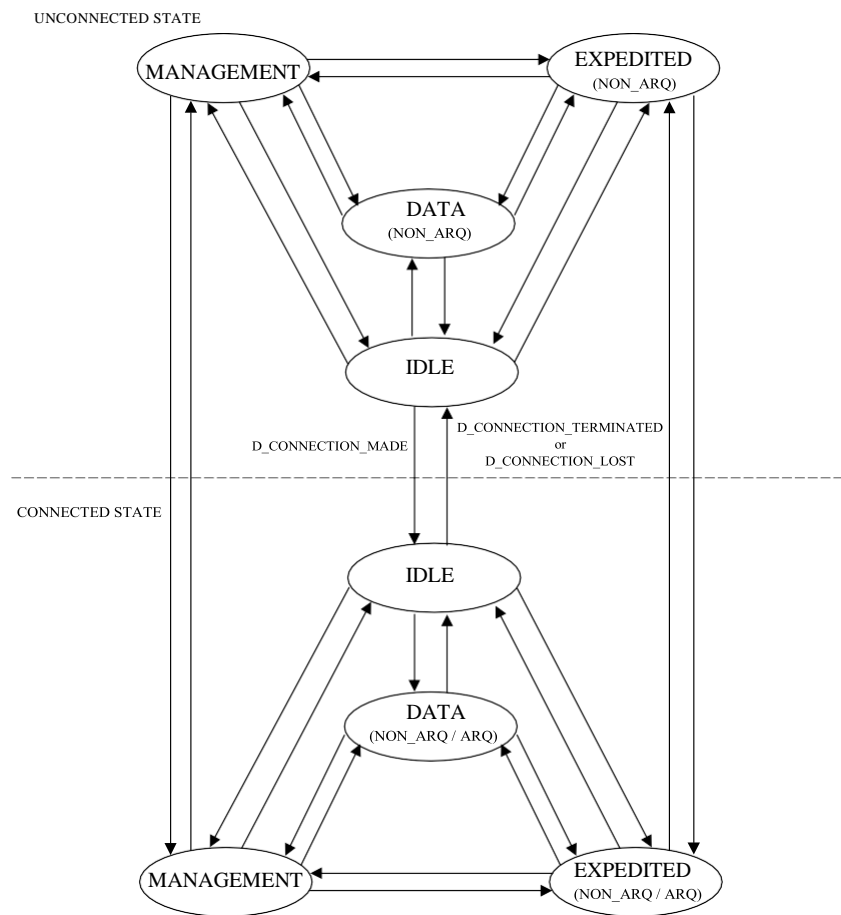
Figure C-54. Nominal State and Transition Diagram (Data Transfer Sublayer)

C.7.1.1. State Transition Rules

The transitions between DTS States and the actions that arise from given events **shall** [1] be as defined in the tables presented in the subsections that follow.

DTS states, transitions, and actions **shall** [2] be defined and maintained with respect to a specified remote node address.

In all of the tables below, "PDU received" refers to a PDU received that is addressed to the node in question from the specified remote node for which the states are maintained. Action and transitions rules **shall** [3] not occur based on PDUs addressed to other nodes or from a node other than the specified remote node for which the states are maintained.

Likewise, "D_Primitive received" refers to a D_Primitive received from the Channel Access sublayer that references the specified remote node. Action and transitions rules **shall** [4] not occur based on D_Primitives that reference nodes other than the specified remote node.

DTS states **shall** [5] be maintained for each specified node for which a connection or ARQ protocol must be

maintained. [Note: If multiple links, as defined by the Channel Access or Subnetwork Interface sublayers, must be maintained simultaneously, then a set of DTS State Variables must be maintained for each of the links. If links are not maintained simultaneously, DTS State variables may be reused for any node]

C.7.1.1.1.   IDLE(UNCONNECTED) State Transition Rules

When in the IDLE(UNCONNECTED) State, the Data Transfer Sublayer **shall** [(1)] respond to reception of a D_Primitive from the Channel Access Sublayer as specified in the following Table:

Table C-9. Event/Transition/Action Rules for IDLE (UNCONNECTED) State: CAS-Related Events

Received D_Primitive/ Event	State Transition to:	Action & Comments
D_CONNECTION_MADE	IDLE(CONNECTED)	INITIALIZE State Variables for ARQ processing, in accordance with Section C.7.2 [Note: CAS has accepted a connection request from the remote node.]
D_CONNECTION_ TERMINATED	IDLE(UNCONNECTED) , i.e., No Change.	IGNORE [Note: this is an anomalous event, the DTS may notify the CAS that an attempt was made to terminate a connection that did not exist.]
D_UNIDATA_REQUEST (ARQ)	IDLE(UNCONNECTED) , i.e., No Change.	REPLY w/ D_UNIDATA_INDICATION (REJECT). [Note: reliable data cannot be sent over a nonexistent connection.]
D_EXPEDITED_UNIDATA_ REQUEST (ARQ)	IDLE(UNCONNECTED) , i.e., No Change.	REPLY w/ D_EXPEDITED_UNIDATA_INDICATION (REJECT). [Note: reliable data cannot be sent over a nonexistent connection.]
D_UNIDATA_REQUEST (NON-ARQ)	DATA(UNCONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant NON-ARQ DATA Type 7 D_PDUs
D_EXPEDITED_UNIDATA_ REQUEST (NON_ARQ)	EXPEDITED DATA (UNCONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant EXPEDITED NON-ARQ DATA Type 8 D_PDUs

When in the IDLE(UNCONNECTED) State, the Data Transfer Sublayer **shall** [(2)] respond to reception of a D_PDU from the lower layers as specified in the following Table:

Table C-10. Event/Transition/Action Rules for

IDLE (UNCONNECTED) State: Reception-Related Events

Received D_PDU Event	State Transition to:	Action & Comments
DATA Type 0/9 D_PDU	IDLE(UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
ACK Type 1/10 D_PDU	IDLE(UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED

DATA-ACK Type 2/11 D_PDU	IDLE(UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
RESET/WIN RESYNC Type 3/12 D_PDU	IDLE(UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
EXPEDITED-DATA Type 4 D_PDU	IDLE(UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
EXPEDITED-ACK Type 5 D_PDU	IDLE(UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
MANAGEMENT Type 6 D_PDU	MANAGEMENT (UNCONNECTED)	PROCESS Type 6 D_PDU; RESPOND if required.
NON-ARQ DATA Type 7 D_PDU	DATA(UNCONNECTED)	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if received complete C_PDU
EXPEDITED NON_ARQ Type 8 D_PDU	EXPEDITED-DATA (UNCONNECTED)	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if the DTS received a complete C_PDU
EXTENSION Type 13 D_PDU (values 1-6)	IDLE(UNCONNECTED) , i.e., No Change.	Contents of D_PDU Handled by MAC Layer
PADDING Type 14 D_PDU	IDLE(UNCONNECTED) , i.e., No Change.	Ignore
WARNING Type 15 D_PDU	IDLE(UNCONNECTED) , i.e., No Change.	CEASE action indicated in warning; NOTIFY CAS using D_WARNING_RECEIVED

C.7.1.1.2. DATA(UNCONNECTED) State Transition Rules

When in the DATA(UNCONNECTED) State, the Data Transfer Sublayer **shall** [1] respond to reception of a D_Primitive from the Channel Access Sublayer as specified in the following Table:

Table C-11. Event/Transition/Action Rules for DATA (UNCONNECTED) State: CAS-Related Events

Received D_Primitive Event	State Transition to:	Action & Comments
D_CONNECTION_MADE	IDLE(CONNECTED)	INITIALIZE State Variables for ARQ processing, in accordance with Section C.7.2 [Note: CAS has accepted a connection request from the remote node.]
D_CONNECTION_ TERMINATED	DATA (UNCONNECTED) , i.e., No Change.	IGNORE [Note: this is an anomalous event, the DTS may notify the CAS that an attempt was made to terminate a connection that did not exist.]
D_UNIDATA_REQUEST (ARQ)	DATA (UNCONNECTED) , i.e., No Change.	REPLY w/ D_UNIDATA_INDICATION (REJECT). [Note: reliable data cannot be sent over a nonexistent connection.]

D_EXPEDITED_UNIDATA_ REQUEST (ARQ)	DATA(UNCONNECTED) , i.e., No Change.	REPLY w/ D_EXPEDITED_UNIDATA_INDICATION (REJECT). [Note: reliable data cannot be sent over a nonexistent connection.]
D_UNIDATA_REQUEST (NON-ARQ)	DATA(UNCONNECTED) , i.e., No Change.	SEGMENT C_PDU; COMPOSE and SEND all resultant NON-ARQ DATA Type 7 D_PDUs
D_EXPEDITED_UNIDATA_ REQUEST (NON_ARQ)	EXPEDITED-DATA (UNCONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant EXPEDITED NON-ARQ DATA Type 8 D_PDUs

When in the DATA(UNCONNECTED) State, the Data Transfer Sublayer **shall** [(2)] respond to reception of a D_PDU from the lower layers as specified in the following Table:

Table C-12. Event/Transition/Action Rules for DATA (UNCONNECTED) State: Reception-Related Events

Received D_PDU Event	State Transition to:	Action & Comments
DATA Type 0/9 D_PDU	DATA(UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
ACK Type 1/10 D_PDU	DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
DATA-ACK Type 2 /11 D_PDU	DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
RESET/WIN RESYNC Type 3/12 D_PDU	DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
EXPEDITED-DATA Type 4 D_PDU	DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
EXPEDITED-ACK Type 5 D_PDU	DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
MANAGEMENT Type 6 D_PDU	MANAGEMENT (UNCONNECTED)	PROCESS Type 6 D_PDU; RESPOND if required.

NON-ARQ DATA Type 7 D_PDU	DATA(UNCONNECTED), i.e., No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if received complete C_PDU
EXPEDITED NON_ARQ Type 8 D_PDU	DATA(UNCONNECTED), i.e., No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_EXPEDITED_UNIDATA_ INDICATION if the DTS received a complete C_PDU
EXTENSION Type 13 D_PDU (values 1-6)	DATA(UNCONNECTED), i.e., No Change.	Contents of D_PDU Handled by MAC Layer
PADDING Type 14 D_PDU	DATA(UNCONNECTED), i.e., No Change.	Ignore
WARNING Type 15 D_PDU	DATA(UNCONNECTED), i.e., No Change.	CEASE action indicated in warning; NOTIFY CAS using D_WARNING_RECEIVED

C.7.1.1.3. EXPEDITED-DATA(UNCONNECTED) State Transition Rules

When in the EXPEDITED-DATA(UNCONNECTED) State, the Data Transfer Sublayer **shall**[1] respond to reception of a D_Primitive from the Channel Access Sublayer as specified in the following Table:

Table C-13. Event/Transition/Action Rules for EXPEDITED-DATA (UNCONNECTED) State: CAS-Related Event

Received D_Primitive Event	State Transition to:	Action & Comments
D_CONNECTION_MADE	IDLE(CONNECTED)	INITIALIZE State Variables for ARQ processing, in accordance with Section C.7.2 [Note: CAS has accepted a connection request from the remote node.]
D_CONNECTION_ TERMINATED	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	IGNORE [Note: this is an anomalous event, the DTS may notify the CAS that an attempt was made to terminate a connection that did not exist.]
D_UNIDATA_REQUEST (ARQ)	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	REPLY w/ D_UNIDATA_INDICATION (REJECT). [Note: reliable data cannot be sent over a nonexistent connection.]
D_EXPEDITED_UNIDATA_ REQUEST (ARQ)	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	REPLY w/ D_EXPEDITED_UNIDATA_INDICATION (REJECT). [Note: reliable data cannot be sent over a nonexistent connection.]
D_UNIDATA_REQUEST (NON-ARQ)	IF completed transmitting all Expedited-Data, change to DATA(UNCONNECTED), OTHERWISE, EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	IF completed sending all Expedited-Data, SEGMENT C_PDU, then, COMPOSE and SEND all resultant NON-ARQ DATA Type 7 D_PDUs; OTHERWISE, QUEUE C_PDU Pending completion of Expedited-Data transmission.

D_EXPEDITED_UNIDATA_ REQUEST (NON_ARQ)	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	SEGMENT C_PDU; COMPOSE and SEND all resultant EXPEDITED NON-ARQ DATA Type 8 D_PDUs

When in the EXPEDITED-DATA(UNCONNECTED) State, the Data Transfer Sublayer **shall**[2] respond to reception of a D_PDU from the lower layers as specified in the following Table:

Table C-14. Event/Transition/Action Rules for EXPEDITED-DATA (UNCONNECTED) State: Reception-Related Event

Received D_PDU/ Event	State Transition to:	Action & Comments
DATA Type 0/9 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
ACK Type 1/10 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
DATA-ACK Type 2 /11D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
RESET/WIN RESYNC Type 3/12 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
EXPEDITED-DATA Type 4 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
EXPEDITED-ACK Type 5 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
MANAGEMENT Type 6 D_PDU	MANAGEMENT (UNCONNECTED)	PROCESS Type 6 D_PDU; RESPOND if required.
NON-ARQ DATA Type 7 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if received complete C_PDU
EXPEDITED NON_ARQ Type 8 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if the DTS received a complete C_PDU
EXTENSION Type 13 D_PDU (values 1-6)	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	Contents of D_PDU Handled by MAC Layer
PADDING Type 14 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	Ignore

WARNING Type 15 D_PDU	EXPEDITED-DATA (UNCONNECTED), i.e., No Change.	CEASE action indicated in warning; NOTIFY CAS using D_WARNING_RECEIVED

C.7.1.1.4. MANAGEMENT(UNCONNECTED) State Transition Rules

When in the MANAGEMENT(UNCONNECTED) State, the Data Transfer Sublayer **shall** [(1)] respond to reception of a D_Primitive from the Channel Access Sublayer as specified in the following Table:

Table C-15. Event/Transition/Action Rules for MANAGEMENT (UNCONNECTED) State: CAS-Related Events

Received D_Primitive/ Event	State Transition to:	Action & Comments
D_CONNECTION_MADE	MANAGEMENT (CONNECTED)	INITIALIZE State Variables for ARQ processing, in accordance with Section C.7.2 [Note: CAS has accepted a connection request from the remote node.]
D_CONNECTION_ TERMINATED	MANAGEMENT (UNCONNECTED) , i.e., No Change.	IGNORE [Note: this is an anomalous event, the DTS may notify the CAS that an attempt was made to terminate a connection that did not exist.]
D_UNIDATA_REQUEST (ARQ)	MANAGEMENT (UNCONNECTED) , i.e., No Change.	REPLY w/ D_UNIDATA_INDICATION (REJECT). [Note: reliable data cannot be sent over a nonexistent connection.]
D_EXPEDITED_UNIDATA_ REQUEST (ARQ)	MANAGEMENT (UNCONNECTED) , i.e., No Change.	REPLY w/ D_EXPEDITED_UNIDATA_INDICATION (REJECT). [Note: reliable data cannot be sent over a nonexistent connection.]
D_UNIDATA_REQUEST (NON-ARQ)	IF Management Protocol completed, DATA(UNCONNECTED), OTHERWISE, MANAGEMENT (UNCONNECTED) , i.e., No Change.	IF Management Protocol completed, SEGMENT C_PDU; COMPOSE and SEND all resultant NON-ARQ DATA Type 7 D_PDUs; OTHERWISE QUEUE C_PDU Pending completion of management protocol.
D_EXPEDITED_UNIDATA_ REQUEST (NON_ARQ)	IF Management Protocol completed, EXPEDITED DATA (UNCONNECTED), OTHERWISE, MANAGEMENT (UNCONNECTED) , i.e., No Change.	IF Management Protocol completed, SEGMENT C_PDU; COMPOSE and SEND all resultant EXPEDITED NON-ARQ DATA Type 8 D_PDUs; OTHERWISE QUEUE C_PDU Pending completion of management protocol.

When in the MANAGEMENT(UNCONNECTED) State, the Data Transfer Sublayer **shall** [(2)] respond to reception of a D_PDU from the lower layers as specified in the following Table:

Table C-16. Event/Transition/Action Rules for MANAGEMENT (UNCONNECTED) State: Reception-Related Events

Received D_PDU Event	State Transition to:	Action & Comments
DATA Type 0/9 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
ACK Type 1/10 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
DATA-ACK Type 2 /11D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
RESET/WIN RESYNC Type 3/12 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
EXPEDITED-DATA Type 4 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
EXPEDITED-ACK Type 5 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =1="Connection-Related D_PDU but Unconnected"; NOTIFY CAS using D_WARNING_TRANSMITTED
MANAGEMENT Type 6 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	PROCESS Type 6 D_PDU; RESPOND if required.
NON-ARQ DATA Type 7 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if received complete C_PDU
EXPEDITED NON_ARQ Type 8 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if the DTS received a complete C_PDU
EXTENSION Type 13 D_PDU (values 1-6)	MANAGEMENT (UNCONNECTED) , i.e., No Change.	Contents of D_PDU Handled by MAC Layer
PADDING Type 14 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	Ignore
WARNING Type 15 D_PDU	MANAGEMENT (UNCONNECTED) , i.e., No Change.	CEASE action indicated in warning; NOTIFY CAS using D_WARNING_RECEIVED

C.7.1.1.5. **IDLE(CONNECTED) State Transition Rules**

When in the IDLE(CONNECTED) State, the Data Transfer Sublayer **shall** [1] respond to reception of a D_Primitive from the Channel Access Sublayer as specified in the following Table:

Table C-17. Event/Transition/Action Rules for IDLE (CONNECTED) State: CAS-Related Event

Received D_Primitive Event	State Transition to:	Action & Comments
D_CONNECTION_MADE	IDLE(CONNECTED), i.e. No Change	IGNORE [Note: this is an anomalous event, the DTS may notify the CAS that an attempt was made to make a connection that already exists.]
D_CONNECTION_ TERMINATED	IDLE(UNCONNECTED)	TERMINATE ARQ processing; RESET State Variables.
D_UNIDATA_REQUEST (ARQ)	DATA(CONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant ARQ DATA Type 0 D_PDUs;
D_EXPEDITED_UNIDATA_ REQUEST (ARQ)	EXPEDITED-DATA(CONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant ARQ DATA Type 4 D_PDUs;
D_UNIDATA_REQUEST (NON-ARQ)	DATA(CONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant NON-ARQ DATA Type 7 D_PDUs
D_EXPEDITED_UNIDATA_ REQUEST (NON_ARQ)	EXPEDITED DATA (CONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant EXPEDITED NON-ARQ DATA Type 8 D_PDUs

When in the IDLE(CONNECTED) State, the Data Transfer Sublayer **shall** [2] respond to reception of a D_PDU from the lower layers as specified in the following Table:

Table C-18. Event/Transition/Action Rules for IDLE (CONNECTED) State: Reception-Related Event

Received D_PDU/ Event	State Transition to:	Action & Comments
DATA Type 0/9 D_PDU	DATA(CONNECTED)	PROCESS the D_PDU; REASSEMBLE C_PDU segments; UPDATE ARQ State Variables; DELIVER received C_PDU to CAS using D_UNIDATA_INDICATION if received complete C_PDU; COMPOSE and SEND ACK Type 1 D_PDU or DATA-ACK Type 2 D_PDU.
ACK Type 1/10 D_PDU	DATA(CONNECTED)	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =3="Invalid D_PDU for Current State"; NOTIFY CAS using D_WARNING_TRANSMITTED

DATA-ACK Type 2/11 D_PDU	DATA(CONNECTED)	PROCESS the D_PDU; REASSEMBLE C_PDU segments; UPDATE ARQ State Variables; DELIVER received C_PDU to CAS using D_UNIDATA_INDICATION if received complete C_PDU; COMPOSE and SEND ACK Type 1 D_PDU or DATA-ACK Type 2 D_PDU as appropriate; as appropriate, NOTIFY CAS using D_UNIDATA_REQUEST_ CONFIRM or D_UNIDATA_REQUEST_ REJECTED.
RESET/WIN RESYNC Type 3/12 D_PDU	IDLE(CONNECTED), i.e. No Change	COMPOSE and SEND RESET/WIN RESYNC Type 3 D_PDU as required; EXECUTE FULL RESET/WIN RESYNC Protocol if required.
EXPEDITED-DATA Type 4 D_PDU	EXPEDITED-DATA (CONNECTED)	PROCESS the D_PDU; REASSEMBLE C_PDU segments; DELIVER to CAS using D_EXPEDITED_UNIDATA_INDICATION if received complete C_PDU; COMPOSE and SEND EXPEDITED ACK Type 5 D_PDU
EXPEDITED-ACK Type 5 D_PDU	EXPEDITED-DATA (CONNECTED)	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =3="Invalid D_PDU for Current State"; NOTIFY CAS using D_WARNING_TRANSMITTED
MANAGEMENT Type 6 D_PDU	MANAGEMENT (CONNECTED)	PROCESS Type 6 D_PDU; RESPOND if required.
NON-ARQ DATA Type 7 D_PDU	DATA(CONNECTED)	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_ INDICATION if received complete C_PDU
EXPEDITED NON_ARQ Type 8 D_PDU	EXPEDITED-DATA (CONNECTED)	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_EXPEDITED_ UNIDATA_INDICATION if a complete C_PDU is received
EXTENSION Type 13 D_PDU (values 1-6)	IDLE(CONNECTED), i.e. No Change	Contents of D_PDU Handled by MAC Layer
PADDING Type 14 D_PDU	IDLE(CONNECTED), i.e. No Change	Ignore
WARNING Type 15 D_PDU	IDLE(CONNECTED), i.e. No Change	CEASE action indicated in warning; NOTIFY CAS using D_WARNING_RECEIVED

C.7.1.1.6. DATA(CONNECTED) State Transition Rules

When in the DATA(CONNECTED) State, the Data Transfer Sublayer **shall** [(1)] respond to reception of a D_Primitive from the Channel Access Sublayer as specified in the following Table:

Table C-19. Event/Transition/Action Rules for DATA (CONNECTED) State: CAS-Related Events

Received D_Primitive Event	State Transition to:	Action & Comments

D_CONNECTION_MADE	DATA(CONNECTED)	IGNORE [Note: this is an anomalous event, the DTS may notify the CAS that an attempt was made to make a connection that already exists.]
D_CONNECTION_ TERMINATED	IDLE(UNCONNECTED)	TERMINATE ARQ processing; RESET State Variables.
D_UNIDATA_REQUEST (ARQ)	DATA(CONNECTED)	SEGMENT C_PDU; UPDATE State ARQ Variables; COMPOSE and SEND all resultant DATA Type 0 D_PDUs;
D_EXPEDITED_UNIDATA_ REQUEST (ARQ)	EXPEDITED- DATA(CONNECTED)	SEGMENT C_PDU; UPDATE State ARQ Variables; COMPOSE and SEND all resultant EXPEDITED-DATA Type 4 D_PDUs;
D_UNIDATA_REQUEST (NON-ARQ)	DATA(CONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant NON-ARQ DATA Type 7 D_PDUs
D_EXPEDITED_UNIDATA_ REQUEST (NON_ARQ)	EXPEDITED DATA (CONNECTED)	SEGMENT C_PDU; COMPOSE and SEND all resultant EXPEDITED NON-ARQ DATA Type 8 D_PDUs

When in the DATA(CONNECTED) State, the Data Transfer Sublayer **shall** [2] respond to reception of a D_PDU from the lower layers as specified in the following Table:

Table C-20. Event/Transition/Action Rules for DATA (CONNECTED) State: Reception-Related Events

Received D_PDU Event	State Transition to:	Action & Comments
DATA Type 0/9 D_PDU	DATA(CONNECTED) , i.e., No Change.	PROCESS the D_PDU; REASSEMBLE C_PDU segments; UPDATE ARQ State Variables; DELIVER received C_PDU to CAS using D_UNIDATA_INDICATION if a complete C_PDU received; COMPOSE and SEND ACK Type 1 D_PDU or DATA-ACK Type 2 D_PDU.
ACK Type 1/10 D_PDU	DATA (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; UPDATE ARQ State Variables; ADVANCE Flow-Control WINDOW; SEND additional DATA Type 0 D_PDUs, if any and as able; as appropriate, NOTIFY CAS using D_UNIDATA_REQUEST_ CONFIRM or D_UNIDATA_REQUEST_ REJECTED.
DATA-ACK Type 2 /11D_PDU	DATA (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; REASSEMBLE C_PDU segments; UPDATE ARQ State Variables; DELIVER received C_PDU to CAS using D_UNIDATA_INDICATION if received complete C_PDU; COMPOSE and SEND ACK Type 1 D_PDU or DATA-ACK Type 2 D_PDU as appropriate; as appropriate, NOTIFY CAS using D_UNIDATA_REQUEST_ CONFIRM or D_UNIDATA_REQUEST_ REJECTED.

RESET/WIN RESYNC Type 3/12 D_PDU	If FULL_RESET_CMD, IDLE(CONNECTED), otherwise, DATA(CONNECTED), i.e. No Change	COMPOSE and SEND RESET/WIN RESYNC Type 3 D_PDU as required; EXECUTE FULL RESET/WIN RESYNC Protocol if required.
EXPEDITED-DATA Type 4 D_PDU	EXPEDITED-DATA (CONNECTED)	PROCESS the D_PDU; REASSEMBLE C_PDU segments; DELIVER to CAS using D_EXPEDITED_UNIDATA_ INDICATION if received complete C_PDU; COMPOSE and SEND EXPEDITED ACK Type 5 D_PDU
EXPEDITED-ACK Type 5 D_PDU	EXPEDITED-DATA (CONNECTED)	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =3="Invalid D_PDU for Current State"; NOTIFY CAS using D_WARNING_TRANSMITTED
MANAGEMENT Type 6 D_PDU	MANAGEMENT (CONNECTED	PROCESS Type 6 D_PDU; RESPOND if required.
NON-ARQ DATA Type 7 D_PDU	DATA(CONNECTED), i.e. No Change	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_ INDICATION if received complete C_PDU
EXPEDITED NON_ARQ Type 8 D_PDU	EXPEDITED-DATA (CONNECTED)	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_EXPEDITED_ UNIDATA_INDICATION if the DTS received a complete C_PDU
EXTENSION Type 13 D_PDU (values 1-6)	DATA(CONNECTED), i.e. No Change	Contents of D_PDU Handled by MAC Layer
PADDING Type 14 D_PDU	DATA(CONNECTED), i.e. No Change	Ignore
WARNING Type 15 D_PDU	DATA(CONNECTED), i.e., No Change.	CEASE action indicated in warning; NOTIFY CAS using D_WARNING_RECEIVED

C.7.1.1.7. EXPEDITED-DATA (CONNECTED) State Transition Rules

When in the EXPEDITED-DATA(CONNECTED) State, the Data Transfer Sublayer **shall** [1] respond to reception of a D_Primitive from the Channel Access Sublayer as specified in the following Table:

Table C-21. Event/Transition/Action Rules for EXPEDITED-DATA (CONNECTED) State: CAS-Related Events

Received D_Primitive Event	State Transition to:	Action & Comments
D_CONNECTION_MADE	EXPEDITED-DATA (CONNECTED), i.e. No Change	IGNORE [Note: this is an anomalous event, the DTS may notify the CAS that an attempt was made to make a connection that already exists.]
D_CONNECTION_ TERMINATED	IDLE(UNCONNECTED)	TERMINATE ARQ processing; RESET State Variables.

D_UNIDATA_REQUEST (ARQ)	in Accordance with Section C.7.1.3: IF completed transmitting all Expedited-Data, change to DATA(CONNECTED), OTHERWISE, EXPEDITED-DATA (CONNECTED), i.e., No Change.	in Accordance with Section C.7.1.3: IF completed sending all Expedited-Data, SEGMENT C_PDU, UPDATE State Variables, then, COMPOSE and SEND all resultant DATA Type 0 D_PDUs; OTHERWISE, QUEUE C_PDU Pending completion of Expedited-Data transmission.
D_EXPEDITED_UNIDATA_ REQUEST (ARQ)	in Accordance with Section C.7.1.3: EXPEDITED-DATA (CONNECTED), i.e. No Change	in Accordance with Section C.7.1.3: SEGMENT C_PDU; UPDATE State ARQ Variables; COMPOSE and SEND all resultant EXPEDITED-DATA Type 4 D_PDUs;
D_UNIDATA_REQUEST (NON-ARQ)	in Accordance with Section C.7.1.3: IF completed transmitting all Expedited-Data, change to DATA(CONNECTED), OTHERWISE, EXPEDITED-DATA (CONNECTED), i.e., No Change.	in Accordance with Section C.7.1.3: IF completed sending all Expedited-Data, SEGMENT C_PDU, then, COMPOSE and SEND all resultant NON-ARQ DATA Type 7 D_PDUs; OTHERWISE, QUEUE C_PDU Pending completion of Expedited-Data transmission.
D_EXPEDITED_UNIDATA_ REQUEST (NON_ARQ)	EXPEDITED-DATA (CONNECTED), i.e. No Change.	SEGMENT C_PDU; COMPOSE and SEND all resultant EXPEDITED NON-ARQ DATA Type 8 D_PDUs

When in the EXPEDITED-DATA(CONNECTED) State, the Data Transfer Sublayer **shall** [(2)] respond to reception of a D_PDU from the lower layers as specified in the following Table:

Table C-22. Event/Transition/Action Rules for EXPEDITED-DATA (CONNECTED) State: Reception-Related Events

Received D_PDU Event	State Transition to:	Action & Comments
DATA Type 0/9 D_PDU	EXPEDITED-DATA (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; REASSEMBLE C_PDU segments; UPDATE ARQ State Variables; DELIVER received C_PDU to CAS using D_UNIDATA_INDICATION if received complete C_PDU; COMPOSE and SEND ACK Type 1 D_PDU
ACK Type 1/10 D_PDU	EXPEDITED-DATA (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; UPDATE ARQ State Variables; ADVANCE Flow-Control WINDOW; as appropriate, NOTIFY CAS using D_UNIDATA_REQUEST_ CONFIRM or D_UNIDATA_REQUEST_ REJECTED.

ANNEX C
STANAG 5066
Edition 4

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

DATA-ACK Type 2 /11D_PDU	EXPEDITED-DATA (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; REASSEMBLE C_PDU segments; UPDATE ARQ State Variables; DELIVER received C_PDU to CAS using D_UNIDATA_INDICATION if received complete C_PDU; COMPOSE and SEND ACK Type 1 D_PDU; as appropriate, NOTIFY CAS using D_UNIDATA_REQUEST_ CONFIRM or D_UNIDATA_REQUEST_ REJECTED.
RESET/WIN RESYNC Type 3/12 D_PDU	If FULL_RESET_CMD, IDLE(CONNECTED), otherwise, DATA(CONNECTED)	COMPOSE and SEND RESET/WIN RESYNC Type 3 D_PDU as required; EXECUTE FULL RESET/WIN RESYNC Protocol if required.
EXPEDITED-DATA Type 4 D_PDU	EXPEDITED-DATA (CONNECTED), i.e. No Change.	PROCESS the D_PDU; REASSEMBLE C_PDU segments; DELIVER to CAS using D_EXPEDITED_UNIDATA_INDICATION if received complete C_PDU; COMPOSE and SEND EXPEDITED ACK Type 5 D_PDU
EXPEDITED-ACK Type 5 D_PDU	EXPEDITED-DATA (CONNECTED), i.e. No Change.	PROCESS the D_PDU; UPDATE ARQ State Variables; ADVANCE Flow-Control WINDOW; SEND additional EXPEDITED-DATA D_PDUs, if any and as able; as appropriate, NOTIFY CAS using D_EXPEDITED_ UNIDATA_REQUEST_ CONFIRM or D_EXPEDITED_UNIDATA_REQUEST_ REJECTED.
MANAGEMENT Type 6 D_PDU	MANAGEMENT (CONNECTED)	PROCESS Type 6 D_PDU; RESPOND if required.
NON-ARQ DATA Type 7 D_PDU	EXPEDITED-DATA (CONNECTED), i.e. No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_ INDICATION if received complete C_PDU
EXPEDITED NON_ARQ Type 8 D_PDU	EXPEDITED-DATA (CONNECTED), i.e. No Change	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_EXPEDITED_ UNIDATA_INDICATION if the DTS received a complete C_PDU
EXTENSION Type 13 D_PDU (values 1-6)	EXPEDITED-DATA (CONNECTED), i.e. No Change	Contents of D_PDU Handled by MAC Layer
PADDING Type 14 D_PDU	EXPEDITED-DATA (CONNECTED), i.e. No Change	Ignore
WARNING Type 15 D_PDU	EXPEDITED-DATA (CONNECTED) , i.e., No Change.	CEASE action indicated in warning; NOTIFY CAS using D_WARNING_RECEIVED

C.7.1.1.8. **MANAGEMENT (CONNECTED) State Transition Rules**

When in the MANAGEMENT(CONNECTED) State, the Data Transfer Sublayer **shall** [1] respond to reception of a D_Primitive from the Channel Access Sublayer as specified in the following Table:

Table C-23. Event/Transition/Action Rules for MANAGEMENT (CONNECTED) State: CAS-Related Events

Received D_Primitive/ Event	State Transition to:	Action & Comments
D_CONNECTION_MADE	MANAGEMENT (CONNECTED) , i.e., No Change.	IGNORE [Note: this is an anomalous event, the DTS may notify the CAS that an attempt was made to make a connection that already exists.]
D_CONNECTION_ TERMINATED	IDLE (UNCONNECTED)	TERMINATE ARQ processing; RESET State Variables.
D_UNIDATA_REQUEST (ARQ)	IF Management Protocol completed, DATA(CONNECTED), OTHERWISE, MANAGEMENT (CONNECTED) , i.e., No Change.	IF Management Protocol completed, SEGMENT C_PDU; UPDATE State ARQ Variables; COMPOSE and SEND all resultant DATA Type 0 D_PDUs; OTHERWISE QUEUE C_PDU Pending completion of management protocol.
D_EXPEDITED_UNIDATA_ REQUEST (ARQ)	IF Management Protocol completed, EXPEDITED DATA (CONNECTED), OTHERWISE, MANAGEMENT (CONNECTED), i.e., No Change.	IF Management Protocol completed, SEGMENT C_PDU; UPDATE State ARQ Variables; COMPOSE and SEND all resultant EXPEDITED-DATA Type 4 D_PDUs; OTHERWISE QUEUE C_PDU Pending completion of management protocol.
D_UNIDATA_REQUEST (NON-ARQ)	IF Management Protocol completed, DATA(CONNECTED), OTHERWISE, MANAGEMENT (CONNECTED) , i.e., No Change.	IF Management Protocol completed, SEGMENT C_PDU; COMPOSE and SEND all resultant NON-ARQ DATA Type 7 D_PDUs; OTHERWISE QUEUE C_PDU Pending completion of management protocol.
D_EXPEDITED_UNIDATA_ REQUEST (NON_ARQ)	IF Management Protocol completed, EXPEDITED DATA (CONNECTED), OTHERWISE, MANAGEMENT (CONNECTED) , i.e., No Change.	IF Management Protocol completed, SEGMENT C_PDU; COMPOSE and SEND all resultant EXPEDITED NON-ARQ DATA Type 8 D_PDUs; OTHERWISE QUEUE C_PDU Pending completion of management protocol.

When in the MANAGEMENT(CONNECTED) State, the Data Transfer Sublayer **shall** [2] respond to reception of a D_PDU from the lower layers as specified in the following Table:

Table C-24. Event/Transition/Action Rules for MANAGEMENT (CONNECTED) State: Reception-Related Events

Received D_PDU Event	State Transition to:	Action & Comments
DATA Type 0/9 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; REASSEMBLE C_PDU segments; UPDATE ARQ State Variables; DELIVER received C_PDU to CAS using D_UNIDATA_INDICATION if received complete C_PDU; COMPOSE ACK Type 1 D_PDU for deferred transmission on transition to an appropriate state.
ACK Type 1/10 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; UPDATE ARQ State Variables; ADVANCE Flow-Control WINDOW; as appropriate, NOTIFY CAS using D_UNIDATA_REQUEST_ CONFIRM or D_UNIDATA_REQUEST_ REJECTED.
DATA-ACK Type 2 /11D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; REASSEMBLE C_PDU segments; UPDATE ARQ State Variables; DELIVER received C_PDU to CAS using D_UNIDATA_INDICATION if received complete C_PDU; COMPOSE ACK Type 1 D_PDU for deferred transmission on transition to an appropriate state; as appropriate, NOTIFY CAS using D_UNIDATA_REQUEST_ CONFIRM or D_UNIDATA_REQUEST_ REJECTED.
RESET/WIN RESYNC Type 3/12 D_PDU	MANAGEMENT (CONNECTED)	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =3="Invalid D_PDU for Current State"; NOTIFY CAS using D_WARNING_TRANSMITTED; CONTINUE w/ MANAGEMENT Protocol
EXPEDITED-DATA Type 4 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	PROCESS the D_PDU; REASSEMBLE C_PDU segments; DELIVER to CAS using D_EXPEDITED_UNIDATA_INDICATION if received complete C_PDU; COMPOSE EXPEDITED ACK Type 5 D_PDU for deferred transmission on transition to an appropriate state.
EXPEDITED-ACK Type 5 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	COMPOSE and SEND WARNING Type 15 D_PDU, w/ Reason =3="Invalid D_PDU for Current State"; NOTIFY CAS using D_WARNING_TRANSMITTED
MANAGEMENT Type 6 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	PROCESS Type 6 D_PDU; RESPOND if required.
NON-ARQ DATA Type 7 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if received complete C_PDU
EXPEDITED NON_ARQ Type 8 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	PROCESS the D_PDU and reassemble C_PDU segments; DELIVER to CAS using D_UNIDATA_INDICATION if the DTS received a complete C_PDU
EXTENSION Type 13 D_PDU (values 1-6)	MANAGEMENT (CONNECTED) , i.e., No Change.	Contents of D_PDU Handled by MAC Layer

PADDING Type 14 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	Ignore
WARNING Type 15 D_PDU	MANAGEMENT (CONNECTED) , i.e., No Change.	CEASE action indicated in warning; NOTIFY CAS using D_WARNING_RECEIVED

C.7.1.2. State Rules for Sending and Receiving D_PDUs

A node **shall** [1] transmit only those D_PDUs which are allowed for its current state as follows:

IDLE(UNCONNECTED)	type 13, 14 or 15 D_PDUs
DATA(UNCONNECTED)	type 7, 13, 14 or 15 D_PDUs
EXPEDITED DATA(UNCONNECTED)	type 8, 13, 14 or 15 D_PDUs
MANAGEMENT(UNCONNECTED)	type 6, 13, 14 or 15 D_PDUs
IDLE(CONNECTED)	type 13, 14 15 D_PDUs
DATA(CONNECTED)	type 0, 1, 2, 3, 7, 8, 9, 10, 11, 12, 13, 14 or 15
EXPEDITED DATA(CONNECTED)	type 1, 4, 5, 8, 10, 13, 14 or 15 D_PDUs
MANAGEMENT(CONNECTED)	type 6, 8, 13, 14 or 15 D_PDUs

A node **shall** [2] receive and process all valid D_PDUs regardless of its current state. Transmission of responses to a received D_PDU may be immediate or deferred, as appropriate for the current state and as specified in Section C.7.1.1.

C.7.1.3. D_PDU Processing Rules: EXPEDITED-DATA (CONNECTED) State

A separate Frame Sequence Number counter **shall** [1] be maintained for the transmission of ARQ Expedited Data, distinct from the counter with the same name used for regular-delivery service with D_PDU types 0/9 and 2/11.

A separate C_PDU ID counter **shall** [2] be maintained for the transmission of ARQ Expedited Data, distinct from the counter with the same name used for regular non-ARQ and expedited non-ARQ delivery services with D_PDU types 7 and 8.

Upon entering the EXPEDITED-DATA (CONNECTED) state, the EXPEDITED-DATA D_PDU Frame Sequence Number counter **shall** [3] be set to the value zero (0).

Starting or restarting another ARQ machine (i.e. establishing a link with a new node or re-establishing a link with a previously connected node) **shall** [4] reset the ARQ machine for the EXPEDITED DATA-state.

The processing of D_PDUs containing Expedited Data **shall** [5] proceed according to a *C_PDU-level* stop-and-wait protocol, as follows:
   a. No new Expedited-C_PDUs **shall** [6] be accepted for service until the last D_PDU of a prior Expedited-C_PDU has been acknowledged.
   b. Each time a D_EXPEDITED_UNIDATA_REQUEST Primitive is accepted for service, the Expedited Data D_PDU Frame Sequence Number counter **shall** [7] be reset to the value zero and the C_PDU ID counter **shall** [8] be incremented modulo-16.

Upon exiting the EXPEDITED DATA(CONNECTED) state to another state, all unsent EXPEDITED-DATA C_PDUs (and portions of C_PDUs) **shall** [9] be discarded.

Similarly at a receiving node, transition from the EXPEDITED DATA(CONNECTED) state to another state **shall** [10] result in the deletion of a partially assembled C_PDU. [Note: The decision to delete partially processed C_PDUs when a transition is made to another data transfer state reflects the primary usage of the expedited data transfer service, i.e. the transfer of short, high priority, upper layer peer-to-peer PDUs that require immediate acknowledgement.]

C.7.2. D_PDU Frame-Sequence Numbers and Flow-Control when in the DATA STATE

Prior to reaching DATA STATE, CAS-1 negotiation will have determined whether to use the family of D_PDUs with 16 bit frame sequence number (Types 9, 10, 11 and 12) or the family of D_PDUs with 8 bit frame sequence number (Types 0, 1, 2 and 3). The logic of DATA state processing is the same for each family, with the different frame sequence number size leading to a different modulo (256 for 8 bit and 65,536 for 16 bit). The logic is described using the generic D_PDU names, which is also the same as the 16 bit frame sequence number D_PDU names.

Examples in this section use both numbers. Note that the two families of D_PDU **shall not** be mixed.

Because frame numbers are operated upon using modulo arithmetic, it is convenient to represent the ARQ sliding-window that controls D_PDU flow as a segment of a circular buffer as shown in Figure C-55 and Figure C-56.

Figure C-55. Sequence-number Space (8 bit): Integers 0..255

Figure C-56. Sequence-number Space (16 bit): Integers 0..65,535

The Data Transfer Sublayer **shall** [(1)] satisfy the following requirements for D_PDU flow control:

a) Each node **shall** [(2)] maintain a transmit and a receive flow-control window buffer for each connection supported.

b) The frame sequence numbers **shall** [(3)] be assigned uniquely and sequentially in an ascending modulo 256 or 65,536 order during the segmentation of the C_PDU into D_PDUs.

c) The frame sequence numbers **shall** [(4)] not be released for reuse with another D_PDU until the receiving node has acknowledged the D_PDU to which the number is assigned.

d) The transmit lower window edge (TX LWE) **shall** [(5)] indicate the lowest-numbered outstanding unacknowledged D_PDU (lowest-numbered allowing for the modulo 256 or 65,536 operations).

e) The transmit upper window edge (TX UWE) **shall** [(6)] be the number of the last new D_PDU that was transmitted (highest D_PDU number, allowing for the modulo 256 or 65,536 operations).

f) The difference (as a modulo-256 or 65,536 arithmetic operator) between the TX UWE and TX LWE – 1 **shall** be equal to the "current transmitter window size".

g) The "maximum window size" **shall** [(8)] equal 128 or 32,768.

h) The "maximum allowable window size" may be a node-configurable parameter (this is recommended) and **shall** [(9)] not exceed the "maximum window size".

i) The "current transmitter window size" at any moment is variable as a function of the current TX UWE and TX LWE and **shall** [10] not exceed the "maximum allowable window size". This allows for no more than 128 or 32,768 ("maximum window size") outstanding unacknowledged D_PDUs in any circumstance.

j) If the "current transmitter window size" equals the "maximum allowable window size", no additional new D_PDUs **shall** [11] be transmitted until the TX LWE has been advanced and the newly computed difference (modulo 256 or 65,536) between the TX UWE and the TX LWE – 1 is less than the maximum allowable window size.

k) The receive lower window edge (RX LWE) **shall** [12] indicate the oldest D_PDU number that has not been received (lowest D_PDU number, allowing for modulo 256 or 65,536 arithmetic operations).

l) The receive lower window edge (RX LWE) **shall** [13] not be decreased when retransmitted D_PDUs are received that are copies of D_PDUs received previously.

m) The receive upper window edge (RX UWE) **shall** [14] be the frame-sequence number of the last new D_PDU received. [Note: More explicitly, the RX UWE is the Frame Sequence Number of the received D_PDU which is the greatest distance (modulo 256 or 65,536) from the RX LWE. The RX UWE increases monotonically as D_PDUs with higher (mod 256 or 65,536) TX Frame Sequence Numbers are received; it does not move back when retransmitted D_PDUs are received.]

n) D_PDUs with TX FSN falling outside the maximum window size of 128 or 32,768 **shall** [14] not be accepted or acknowledged by the receiver node.

o) If the value of the RX LWE field in any ACK Type 1 or DATA-ACK Type 2 D_PDU is greater than the TX LWE, the Data Transfer Sublayer **shall** [16] declare all D_PDUs with Frame Sequence Numbers between the TX LWE and the RX LWE value as acknowledged, and **shall** [17] advance the TX LWE by setting it equal to the value of the RX LWE.

p) The initial condition of the window edge pointers (e.g., on the initialization of a new link) **shall** [18] be as follows:
- TX LWE = 0
- TX UWE = 255 or 65,535
- RX LWE = 0
- RX UWE = 255 or 65,535

[Note that, although the flow control limitations limit the number of D_PDUs to be transmitted in one transmission, other protocol parameters will also effect this in an indirect way; the structure of the EOT parameter allows a maximum transmission interval of about 2 minutes. If a D_PDU size of 200 bytes is used at 75 bps, only 5 D_PDUs can be transmitted. However, even if a node receives no acknowledgements, it would be possible to transmit many more D_PDUs before transmission would halt due to flow control restrictions. At higher speeds, the 128 D_PDU limit of 8bit frame sequence numbers impacts performance, and so use of 16 bit frame sequence number is desirable.]

The nominal relationships between the ARQ Flow-Control variables specified above are shown in Figure C-57 Figure C-58.
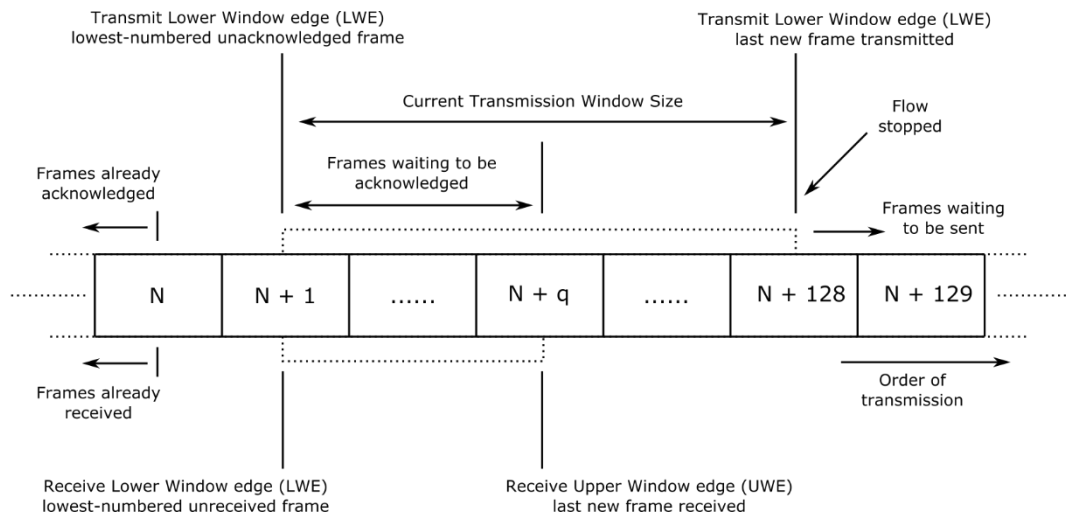
ARQ Flow-Control Variables at the Receiver

Figure C-57. ARQ Flow Control Window Variables (8 bit)

ARQ Flow-Control Variables at the Receiver

Figure C-58. ARQ Flow Control Window Variables (16 bit)

C.7.3. Synchronisation of the ARQ Machine

Procedures for synchronizing the ARQ machine for a link are specified in the following paragraphs.

C.7.3.1. Initial Synchronisation

On starting an ARQ machine (i.e. establishing a link with a new node, or establishing a new link with a previously connected node), the transmit and receive ARQ window-edge pointers **shall** [1] be set to the initial values (as defined in Section C.7.2).

On an HF circuit, it may be desirable to revive a data state connection that has, for example, timed out partway through a file transmission due to a temporary worsening of the propagation conditions. In this case, the loss of data will be minimised if the ARQ machine associated with the connection is re-activated (i.e. the transmit and receive ARQ window-edge pointers are not re-initialised on re-establishing the link). However, this assumption may, for various reasons, *not* be valid (for example, because one of the nodes has experienced a power failure before re-activation) and so a synchronisation verification procedure **shall** [2] be executed whenever a link is re-established.

C.7.3.2. Verification and Maintenance of Synchronisation

The following synchronisation verification procedure **shall** [1] be used to verify on an *ongoing basis* if the peer ARQ processes are in synchronisation and, if required, to effect a reset or re-synchronisation of the peer ARQ window pointers.[3]

Figure C-59 illustrates the fact that, at each end of the node, there is one transmit and one receive window. The dotted lines show which pairs of windows need to remain in synchronisation.

Figure C-59. Synchronization Points

In the event that it is determined that the nodes are not synchronized, the nodes **shall** be synchronized following the procedure set out in Section C.7.4.

C.7.3.2.1. Synchronisation Tests Performed at the Destination Node

Synchronisation tests performed at the destination node make use of the TX lower window edge (LWE) and

TX upper window edge (UWE) flags in conjunction with the TX FRAME SEQ # contained in the header of DATA-ONLY and DATA-ACK frames. The appropriate flag is raised (value = 1) when the TX FRAME SEQ # corresponds to the originating node's transmit ARQ LWE or UWE pointers. The following tests **shall** [2] be used to detect *loss of synchronisation*.

Test 1 : Transmit Upper Window Edge

The purpose of this test is to ensure that the TX UWE of the originating node is within the valid range defined by the destination node window edge pointers. This test **shall** [3] be carried out whenever a D_PDU is received with its TX UWE flag set. If the TX UWE passes this test then the two nodes are *in synchronisation*.

Equation 1 :

$$\text{IN SYNC} = (\text{TX UWE} >= \text{RX UWE}) \textbf{ AND } (\text{TX UWE} <= \text{MAX WIN SIZE} -1 + \text{RX LWE})$$

If the peer ARQ machines are properly synchronised, the TX UWE cannot be less than the RX UWE (as this would indicate that a D_PDU has been received which has not been transmitted). This condition is expressed by the first part of equation 1. It also cannot exceed the limits defined by the maximum window size of 128 or 32,768 and therefore cannot be greater than the {RX LWE + MAX WIN SIZE –1} (modulo 256 or 65,536).   This condition is expressed by the second part of equation 1. Equation 1 can therefore be used to establish whether the TX UWE is *in synchronisation* with the RX window edges and *loss of synchronisation* has occurred if this equation is not satisfied.

> Note 1: Verification of synchronisation on an ongoing basis and, if required, re-synchronisation is the responsibility of the Data Transfer Sublayer. However, under some circumstances a reset or re-synchronisation may be initiated by the Management Sublayer, e.g. following the (re)establishment of a link and as part of some link maintenance procedures.

> Note 2: All the synchronisation tests assume a fixed window size equivalent to the maximum window size of 128 or 32,768 frames. Although the STANAG permits the use of a variable transmit window size, this information is not transmitted over the air. The destination node therefore has no knowledge of the window size of the originating node and so cannot take it into account in the synchronisation tests. The tests should still be applied when the window size is varied but in this case some out-of-synchronisation conditions will not be detected

The *in synchronisation* and *out of synchronisation* regions of the FSN circle described by equation 1 are illustrated graphically in Figure C-60.

Figure C-60. Showing IN-SYNC and OUT-OF-SYNC Regions of the FSN Circle (Equation 1)

<u>Test 2 : Transmit Lower Window Edge</u>

The purpose of this test is to ensure that the TX LWE of the originating node is within the valid range defined by the destination node window edge pointers. This test **shall** [4] be carried out whenever a D_PDU is received with its TX LWE flag set. If the TX LWE passes this test then the two nodes are *in synchronisation*.

Equation 2 :

IN SYNC = (TX LWE >= RX UWE – (MAX WIN SIZE –1)) **AND** (TX LWE <= RX LWE)

If the peer ARQ machines are properly synchronised, the TX LWE should not be outside the bounds defined by the maximum window size as seen from the perspective of the destination node. The TX LWE indicated by the incoming D_PDU therefore cannot be less than the {RX UWE – (MAX WIN SIZE – 1)} (modulo 256). This condition is expressed by the first part of equation 2. The TX LWE also cannot be greater than the RX LWE (as this would indicate that the originating node has marked as acknowledged a frame that the destination node has not yet received). This condition is expressed by the second part of equation 2. Equation 2 can therefore be used to establish whether the TX LWE is *in synchronisation* with the RX window edges and l*oss of synchronisation* has occurred if this equation is not satisfied.

The *in synchronisation* and *out of synchronisation* regions of the FSN circle described by equation 2 are illustrated graphically in Figure C-61.

Figure C-61. Showing IN-SYNC and OUT-OF-SYNC Regions of the FSN Circle (Equation 2)

Test 3 : All Received Frames

The purpose of this test is to ensure that the frame sequence number of a frame received from the originating node is within the valid range defined by the destination node window edge pointers. This test **shall** [5] be carried out on every DATA and DATA-ACK frame received. If the frame sequence number of the incoming frame passes this test then the two nodes are *in synchronisation*.

Equation 3 :

$$\text{IN SYNC} = (\text{TX FSN} <= \text{RX LWE} + (\text{MAX WIN SIZE} - 1)) \textbf{ AND}$$
$$(\text{TX FSN} >= \text{RX UWE} - (\text{MAX WIN SIZE} - 1))$$

If either part of equation 3 is not true, then the frame sequence number falls outside the range defined by the maximum window size of 128 frames and *loss of synchronisation* between the two nodes has occurred.

The *in synchronisation* and *out of synchronisation* regions defined by equation 3 are illustrated graphically in Figure C-62.

Figure C-62. Showing IN-SYNC and OUT-OF-SYNC Regions of the FSN Circle (Equation 3)

C.7.3.2.2. **Synchronisation Tests Performed at the Originating Node**

The purpose of tests carried out at the originating node is to ensure that the frame sequence numbers of acknowledged frames are within the range defined by the transmit window edge pointers. These tests **shall** [6] be applied whenever a DATA-ACK or ACK-ONLY frame is received.

Test 4 : Receive Lower Window Edge

The value of the RX LWE is included in all DATA-ACK and ACK-ONLY frames. The following test is used to determine whether the RX LWE is within the range defined by the TX window edge pointers. If the RX LWE passes this test then the two nodes are *in synchronisation*.

Equation 4 :

$$IN\ SYNC = (RX\ LWE >= TX\ LWE)\ \textbf{AND}\ (RX\ LWE <= TX\ UWE +1)$$

If equation 4 is not satisfied then *loss of synchronisation* has occurred.

The *in synchronisation* and *out of synchronisation* regions defined by equation 4 are illustrated graphically in Figure C-63.

Figure C-63. Showing IN-SYNC and OUT-OF-SYNC Regions of the FSN Circle (Equation 4)

Test 5 : Explicitly Acknowledged Frames

Individual frames may be acknowledged in a DATA-ACK or ACK-ONLY frame by setting bits in the selective ack header field. The frame sequence numbers (FSNs) corresponding to such bits must also fall within the range defined by the transmit window edges if the two nodes are in synchronisation. The following test **shall** [7] be used to determine whether acknowledged FSNs fall within the correct range.

Equation 5 :

$$\text{IN SYNC} = (\text{Acknowledged FSN} > \text{TX LWE}) \textbf{ AND } (\text{Acknowledged FSN} <= \text{TX UWE})$$

If acknowledged FSNs do not satisfy the condition defined in equation 5 then *loss of synchronisation* has occurred.

The *in synchronisation* and *out of synchronisation* regions of the frame sequence number circle of the originating node are illustrated graphically in Figure C-64.

Figure C-64. Showing IN-SYNC and OUT-OF-SYNC Regions of the FSN Circle (Equation 5)

C.7.4. Procedure for use of Type 3 RESET/WIN RESYNC D_PDU

The Type 3 RESET/WIN RESYNC D_PDU, specified in Section C.4.7, supports a number of different resynchronization functions. The procedure for the FULL RESET function is specified in this section.

A FULL RESET procedure **shall** [(1)] initiated by a node sending a Type 3 RESET/WIN RESYNC D_PDU with field values as follows:

- The FULL RESET CMD flag **shall** [(2)] be set to 1;
- The RESET FRAME ID NUMBER **shall** [(3)] be selected from RESET FRAME ID NUMBER

sequence;
- The NEW RECEIVE LWE field **shall** [3] be reset to zero;
- The RESET TX WIN RQST flag **shall** [4] be reset to zero;
- The RESET RX WIN CMD flag **shall** [5] be reset to zero.

The Type 3 D_PDU described immediately above is defined to be a FULL RESET CMD D_PDU. A node receiving a FULL RESET CMD D_PDU **shall** [6] proceed as follows:
- The node **shall** [7] set the transmit and receive window pointers to initial values (as defined in Section C.7.2);
- The node **shall** [8] discard from its transmit queue any partially completed C_PDUs;
- The node **shall** [9] flush its receive buffers;

- The node **shall** [10] respond to the originator of the FULL-RESET-CMD D_PDU by sending it a RESET/WIN RESYNC (Type 3) D_PDU with field values set as follows:

  - The RESET ACK flag **shall** [11] be set to 1;
  - The NEW RECEIVE LWE and RESET FRAME ID NUMBER fields **shall** [12] be reset to zero;
  - The RESET TX WIN RQST, FULL RESET CMD and RESET RX WIN CMD flags **shall** [13] be reset to zero.

The D_PDU described immediately above is defined to be a FULL-RESET-ACK D_PDU. The FULL RESET ACK D_PDU **shall** [14] be sent only in response to the FULL RESET CMD D_PDU.

On receiving the FULL-RESET-ACK D_PDU, the node initiating the FULL RESET procedure **shall** [15] proceed as follows:

- The node **shall** [16] set the transmit and receive window pointers to initial values (as defined in C.6.2)
- The node **shall** [17] discard from its transmit queue any partially completed C_PDUs
- The node **shall** [18] flush its receive buffers

This concludes the FULL RESET procedure.

C.7.5.   Transmitting and Receiving D_PDUs

The DTS state machine and procedures for operation define operation in terms of D_PDUs sent and received. This section describes how this state machine interacts with the channel and how transmissions are structured. It considers three types of configuration:

1. Single Channel, which will send and receive D_PDUs.
2. Duplex Channel Pair.
3. Broadcast Channel.

C.7.5.1. Receiving D_PDUs

Incoming D_PDUs are handled as they are read from the channel. If a D_PDU header is parsed and the header checksum is correct, the following actions are taken:

1. If WTRP (Annex L) is being used:
    a. Pass EOW 13 or 14 to the WTRP layer.
    b. Pass sender and receiver information to the WTRP layer, in support of WTRP promiscuous mode, along with any information on receive signal SNR and quality.
2. Pass any EOW, apart from EOW 13 or 14 to the state machine.
3. Handle EOTs as described below.
4. Make available any information on receive signal SNR and quality to the subsystem described in Section C.7.6.

After this initial processing, validate the Data Checksum for D_PDUs with such a checksum. If this validation fails, make DATA-ONLY and DATA-ACK D_PDUs to the state machine for processing the Non-ARQ with Errors service. Otherwise discard the D_PDU.

Pass Extended D_PDUs of types 1-6 to the WTRP layer. Discard Extended D_PDUs of any other type. All other D_PDUs are passed to the state machine.

C.7.5.1.1.   **Switching To Transmission**

A node operating in single channel mode **may** transmit after receiving a transmission. The rules for when a node may transmit are set out in Annex K (CSMA) and Annex L (WTRP). One or other of these procedures **shall** be followed.

In addition to this, the receiving system **shall** process received EOTs. Any EOT received allows the receiving system to calculate when the transmission will end. Information from multiple EOTs **shall** be used to determine this end point time as accurately as possible. The node **shall not** transmit before this end point time is reached, even if no signal is being received from the channel.

When a transmission is being received and no EOTs are parsed, it is also desirable to pass modem information on the end of transmission to the MAC layer, to help with transmission control.

C.7.5.2. D_PDU Queue

On transmission side, the state machine will lead to a sequence of D_PDUs being generated for transmission. These D_PDUs can be considered to be in a queue. This queue is specified to enable description of the approach to be taken. An implementation is not required to work in this way, provided that the external effect resulting is the same as one from the approach specified here.

Once a D_PDU has been transmitted, it is removed from the queue, with the exception of Non-ARQ-DATA D_PDUs, which **shall** be retained in the queue until the required number of transmissions has been made.

C.7.5.2.1.   **Transmission Order Considerations**

There is no requirement to transmit D_PDUs from the queue in the order that they arrive. A number of

considerations for this choice are noted in the following sections. Two general considerations **shall** apply in all circumstances:

1. When an ARQ C_PDU is segmented, the D_PDUs with the C_PDU fragments **shall** be assigned adjacent frame sequence numbers and initial transmission **shall** be in frame sequence number order without insertion of other D_PDUs, other than retransmissions.
2. Where D_PDUs have differing priority, the first transmission of D_PDUs of higher priority **shall** always be made before the first transmission of D_PDUs of lower priority, when this does not conflict with the previous requirement.
3. When it is determined from acknowledgements that transfer of a DATA-ONLY D_PDU has failed, retransmission of this D_PDU **shall** be given priority over initial transmission of data D_PDUs of the same priority.
4. Where there are not other factors, transmission of D_PDUs **shall** be "oldest first".

Note that the DTS will handle D_PDUs for multiple peers. These rules apply across all peers, as the DTS can transmit data to any valid current destination.

C.7.5.2.2.   Handling TTL Expiry

While D_PDUs are being transmitted, the TTL of one or more C_PDUs in transit **may** expire. There are various reasons this may happen, including handling of higher priority data or delays in transmission due to poor channel conditions. C_PDUs with TTL expired **shall not** be transmitted. One of the following two processes **shall** be followed for ARQ data in order to drop the D_PDUs while maintaining synchronization of the window.

When a C_PDU is dropped due to TTL, it **shall** be rejected to the CAS using either D_UNIDATA_REQUEST_REJECTED or D_EXPEDITED_UNIDATA_REQUEST_REJECTED with reason TTL expired.

A C_PDU **may** be dropped by use of the Drop PDU on each segment, as specified in Section C.4.4. No data is transmitted with this D_PDU, as it is simply used to transfer the segment. The dropped D_PDU **shall** be acknowledged. All segments of the C_PDU **shall** be transmitted and acknowledged.

Alternatively, the window **may** be reset using the procedure described in Section C.7.4. This approach may be more efficient if there are a large number of whole or partial C_PDUs to be dropped.

For non-ARQ data, and queued D_PDUs from C_PDUs with expired TTL **shall** be dropped and not transmitted.

C.7.5.2.1.   Handling ARQ Link Termination

An ARQ link **may** be terminated either because D_CONNECTION_TERMINATED received from the CAS or due to timeouts or other DTS failures. When this happens, C_PDUs that have not been fully processed **shall** be rejected to the CAS using either D_UNIDATA_REQUEST_REJECTED or D_EXPEDITED_UNIDATA_REQUEST_REJECTED with the reason taken from D_CONNECTION_TERMINATED or the local reason.

C.7.5.2.1.   Handling D_PDUs Not Addressed to Local Node

The DTS state machine handles D_PDUs addressed to the local node by unicast or multicast address.

Although other D_PDUs are not handled by the state machine, implementation experience suggests that it is useful to monitor them. The choice and mechanisms to achieve this are an implementation decision.

C.7.5.3. Transmission Length and Parameters - General Considerations

There are a number of parameters that need to be chosen at start of transmission:

1. Transmission Length.
2. Transmission Speed.
3. Interleaver
4. Maximum C_PDU Segment Size.

General guidance and approach on selecting 2-4 is given in Section C.7.6.   Selection specific to the various models are given below.

C.7.5.3.1.   **Modem Block Alignment**

This section defines a procedure only available for Edition 4 (or subsequent) peers. It **shall not** be used when transmitting to an Edition 3 peer.

When severe errors occur on a HF link, they can lead to corruption of a modem data block where no D_PDUs transmitted in that block are received. Because of this, it is desirable to avoid splitting D_PDUs between modem blocks.   DATA-ONLY D_PDUs will often contain C_PDU fragments. The C_PDPU segment size can be adjusted so that D_PDU and modem block boundaries are aligned.

When there are no alignment issues, the optimum choice for C_PDU segment size is one that is equal for all the C_PDU fragments of a given C_PDU, and as large as possible while being less that the maximum C_PDU segment size.

C.7.5.4. Single Channel

In the single channel model, nodes transmit in turn with rules for transmission governed by the chosen MAC layer (WTRP or CSMA).   Each node will make a transmission of up to 127.5 seconds containing some or all of the queued D_PDUs.

C.7.5.4.1.   **Transmission Length Choice**

For single channel, all of the parameters set out in Section C.7.5.3 need to be set at the start of transmission. In particular, the length of transmission needs to be fixed.   The choice of transmission length is an implementation choice.   This section notes a number of considerations in this choice:

1. The maximum transmission length allowed is 127.5 seconds.
2. Use of longer transmissions will optimize throughput.
3. At lower speeds, use of the longest transmission time allowed is generally desirable.
4. At higher speeds is may be desirable to use shorter maximum transmissions. This will reduce throughput and improve latency. This choice may be sensible where a mixture of applications with different QoS requirements are being supported.
5. Note that transmission of a maximum length DATA-ONLY D_PDU at 75 bps takes around 110 seconds, Any constraints on maximum transmission time **shall not** be implemented in a manner that can cause transfer or large D_PDUs at low speeds to be blocked.
6. Interleaver choice is discussed in Section C.7.6. Minimum transmission length will be a single interleaver block.
7. When there is queued traffic of different priorities, a node **may** choose to send only the higher priority traffic, leaving lower priority traffic for a subsequent transmission.

Note that although the length of transmission needs to be fixed at start of transmission, the transmission content does not. It is **recommended** to choose which D_PDUs to transmit as late as possible, so that higher priority D_PDUs or D_PDUs from applications with low latency QoS requirements that arrive after the start of transmission can be included in the transmission.

C.7.5.4.2. **Repeating and Distributing D_PDUs**

It is often beneficial to repeat D_PDUs. In particular:

1. It is generally desirable to repeat ACK-ONLY D_PDUs several times. These are small D_PDUs, and loss of Acks will lead to delays and the overhead of repeat transmission. Repeating transmission will minimize loss.
2. DATA-ONLY D_PDUs from applications which have low latency QoS, particularly when a long transmission is being used with aggressive speed choice optimizing for throughput. Repeating these D_PDUs will improve latency for the application concerned.
3. Non-ARQ D_PDUs with request for repeat transmissions. These repeats **may** be made within a single transmission or over multiple transmissions.

When a D_PDU is repeated, it is desirable to spread the repeated D_PDUs over the transmission. This separation will help to avoid long fades in the channel.

Transmissions will be an exact number of modem block lengths, and not an arbitrary size. This can lead to a situation where there is "extra space" to be used in a transmission. This space **shall** be filled with D_PDUs to the maximum extent possible. Two choices to do this are noted:

1. ACK-ONLY D_PDUs are generally small and useful to repeat. In many situations, they give a practical way to fill space.
2. PADDING D_PDUs are very small. They provide a way to fill space and repeat EOT and EOW transmission.
3. (Edition 4 only) When C_PDUs are being fragmented into multiple D_PDUs, the C_PDU segment size can be carefully chosen to ensure that blocks are filled.

C.7.5.4.3.   EOT Handling

For single channel operations the EOT field **shall** be set in every D_PDU.

A node **shall not** make a transmission when the EOTs from transmission by another node indicate that it is transmitting.

Once an EOT is sent, the EOT in each subsequent D_PDU in that transmission **shall** contain a consistent calculation of the EOT, that is, monotonically decreasing in half-second (0.5 second) intervals.

Calculations of the EOT by a transmitting node **shall** [5] be rounded up to the nearest half-second interval. [Note: rounding-up the calculation of the EOT, rather than truncating it, ensures that the transmission will be completed by the time declared in the EOT field of the D_PDU.]

A node **shall** stop transmitting when the EOT becomes zero. Rules for subsequent behavior are governed by the MAC layer, with CSMA procedures specified in Annex K and WTRP procedures specified in Annex L.

C.7.5.4.4.   Transmitting to Multiple Peers

A node transmitting on a single channel **may** address all traffic to a single peer node. It **may** also address ARQ traffic to multiple nodes, with a CAS-1 soft link maintained with each node.   It **may** also send non-ARQ traffic to unicast or broadcast addresses.   Peers will take it in turns to transmit, with CSMA procedures specified in Annex K and WTRP procedures specified in Annex L.

In determining transmission parameters, the needs of all peers **shall** be taken into consideration.

C.7.5.5. Duplex

A peer can be configured with separate transmit and receive channels.   When this is done, the CAS will ensure that all traffic to the DTS is for that single peer. So when operating over duplex, the DTS will only be handling D_PDUs to be exchanged with a single peer.

C.7.5.5.1.   Transmission Gaps vs Continuous

When a duplex channel has no current data to transmit, two strategies are available:

1. "Gaps".   When there is no more data to transmit, terminate the channel transmission.   This has the advantage of enabling new parameters to be selected for the next transmission, but there is a delay in establishing the new transmission.
2. "Continuous".   Where the transmission is continued, using either ACK-ONLY D_PDUs noting the current window position or PADDING D_PDUs.   This allows the channel to be more responsive to new D_PDUs arriving.

On reception, both modes **shall** be supported.   For transmission, an implementation can choose to support one or both strategies.

C.7.5.5.2.   Transmission Length and Parameters

With a duplex channel, the length of transmission does not need to be decided in advance and transmission can be of arbitrary length.

Transmission speed and interleaver need to be fixed at start of each transmission, using considerations discussed in Section C.7.6.   If conditions and considerations change, a transmission can be terminated and a new transmission started with different parameters.

Sometimes channels will have low volumes of data. For example:

1. XMPP Chat communication with short messages.
2. A reverse channel to one handling a bulk data stream that is just passing back acknowledgements.

For traffic of this nature, it is desirable to choose a conservative transmission that minimizes data loss.

C.7.5.5.3.   **Acknowledgement Handling**

When operating over a duplex channel, acknowledgements can be sent immediately. This contrasts to single channel, where acknowledgements are only sent after a (potentially long) transmission by the peer.

At lower speeds, where transmission of DATA-ONLY D_PDUs can take several seconds or more, it makes sense to send an ACK-ONLY D_PDU for each DATA-ONLY D_PDU received, and it **may** be beneficial to repeat transmission of each ack.   At higher speeds, with many DATA-ONLY D_PDUs arriving each second, it **may** be preferable to acknowledge several DATA-ONLY D_PDUs with a single ACK-ONLY D_PDU.

When data is flowing in both directions, it will generally be desirable to perform acknowledgement with DATA-ACK D_PDUs.   Repeat acknowledgements should use ACK-ONLY D_PDUs.

C.7.5.5.4.   **EOT Handling**

In duplex mode, the length of transmission is generally not known.   When this is the case, the EOT field **shall** be filled with all zeros to communicate that no EOT is set. If the user of a duplex channel determines when a transmission will end, it **may** continue to set EOT to all zeros or it **may** use EOT to record time remaining following the rules of Section C.7.5.4.3.

C.7.5.6. D_PDU Broadcast

A node operating in broadcast mode will send (non-ARQ) data over the channel to multiple nodes and will never receive data back over the channel.

Handling a broadcast channel is similar to a duplex channel and the following rules apply:

1. A broadcast channel has transmission choices of "continuous" and "gaps" described in Section C.7.5.5.1.
2. Transmission length selection follows Section C.7.5.5.2.
3. EOT handling follows Section C.7.5.5.4.

C.7.6. Data Rate Selection

This section describes selection of transmission speed and related parameters.

C.7.6.1. Use of Edition 3 Data Rate Change for Non-Autobaud Waveforms

This specification supports data rate selection only for autobaud waveforms such as STANAG 4539 and STANAG 5069.

STANAG 5066 Edition 3 specifies a mechanism to perform stop/start Data Rate Change with non-autobaud waveforms. An implementation of the current version of Annex C **may** support this Data Rate Change mechanism.

C.7.6.2. Model for Data Rate Selection

The DTS **may** be operated at fixed speed with autobaud or non-autobaud waveforms.

If transmission speed is varied an autobaud waveforms such as STANAG 4539 and STANAG 5069 **shall** be used.

Because the receiver can adapt to transmission speed and interleaver selected, the model is that for each transmission the sender chooses the parameters to use. It can use a number of factors to make this choice including:

1. Receiver Recommendation. This is an important option, as the receiver can generally measure conditions and make a more effective analysis and choice than the sender can using SNR and other measurements. Mechanisms for communicating this recommendation are specified in Section C.7.6.3.

2. Local conditions, including SNR from recent receptions and SNR variation. This can be helpful, noting that sender conditions can be significantly different to receiver.

3. Frame Error Rate from recent transmissions to the peer.

4. Communication results of communication with other peers.

5. QoS requirements of applications. This may be inferred (e.g., that large volumes of data to be transferred requires optimization for throughput and small volumes (likely to be application acks and XMPP chat type applications) should be optimized for latency. This may be known, based on QoS parameters associated with each SAP. The primary choice is to optimize for throughput or latency. Optimizing for throughput will typically mean choosing an aggressive speed with long interleavers and accepting a high frame error rate. Optimizing for latency will typically mean choosing a conservative speed and perhaps repeating D_PDUs.

The algorithms and choices made are left to the implementation. At a minimum, using receiver recommendations and basic QoS analysis is **recommended**.

C.7.6.2.1. Bandwidth Choice for STANAG 5069

STANAG 5069 wideband HF waveform requires specification of a bandwidth (3kHz – 48kHz). Two approaches are supported by this standard:

1. Fixed Bandwidth. The bandwidth is configured to a fixed value for a given peer or the whole network.

2. Use of 4G ALE as specified in MIL-STD-188-141D. The STANAG 5069 design is intended for use with 4G ALE and this is the approach that is anticipated will generally used. Use of STANAG 5069 generally needs to adapt bandwidth to conditions.

C.7.6.3. Receiver Recommendations

The receiver of a transmission is much better placed than the transmitter to determine best transmission speed and related settings. This is because is has access to SNR and other signal-related information on the transmissions received from a sender.

This specification defines a number of EOW messages that the receiver uses to communicate its recommendations to the sender. Because EOWs are part of the standard D_PDU header, EOWs can be sent by a receiver without incurring any protocol overhead.

When a node transmits receiver recommendations, it may send recommendations to multiple nodes in one transmission. These EOW **shall** only be used in D_PDUs sent to a unicast address and the recommendation applies to transmissions from the node identified by that unicast address.

C.7.6.3.1. Standard Operation

The EOWs described in this section **shall not** be transmitted to Edition 3 peers.

Three of the five EOWs are simple communication of receiver recommendations to the sender.

1. MAX-SPEED (TYPE 8) EOW is specified in Section C.6.5. It specifies the recommended transmission speed to achieve maximum throughput and the minimum length of interleaver to be used.

2. LOW-SPEED (TYPE 9) EOW is specified in Section C.6.6. It specifies the maximum recommended speed for low latency traffic and the minimum length of interleaver to be used

3. MAX-SEGMENT (TYPE 10) EOW is specified in Section C.6.7. It specifies a maximum recommended C_PDU Segment Size. This is particularly useful for fixed speed networks, where reducing the size of D_PDUs transmitted is the only option to adapt to poorer conditions.

These EOWs **may** be used by a receiving node to communicate information to the sender to facilitate better speed selection.

SPEED-USED (TYPE 12) EOW is specified in Section C.6.9. SPEED-USED EOW is used to communication transmission speed and interleaver from a sender to a receiver. It is useful when a receiver cannot determine this information locally. The information in this EOW **shall** apply to the current transmission only. It is recommended to only send this EOW when it is determined that the receiver needs it.

SENDER-APPROACH (TYPE 11) EOW is specified in Section C.6.8. SENDER-APPROACH EOW is used to communicate node capabilities to a peer to which it is sending data. There are three pieces of information conveyed:

1. Information as to whether transmission is fixed or variable speed.
2. The strategy the sender is currently using.
3. Whether the node can determine modem transmission speed, to communicate the need for the peer to send SENDER-SPEED EOW.

The first two elements of SENDER-APPROACH EOW facilitate the receiver to prioritize which EOWs to send. Prioritization may mean only sending one type, or repeating transmission of some EOW types more. For example if bulk data is being transferred, the MAX-SPEED EOW is likely to be the most important EOW. Use of SENDER-APPROACH EOW will be most important at low speeds or for short transmissions where it may only be possible to send a limited number of EOWs. At fixed speed, only the recommended C_PDU Segment Size is useful. When choosing the settings, the sender should consider which EOWs it is most interested to receive.

Table C-25 sets out the recommended receiver interpretation of the SENDER-APPROACH EOW value. Priority indicates choice of use when slots are limited and for repeats.

Approach Setting	MAX-SPEED EOW	LOW SPEED EOW	MAX-SEGMENT EOW
Fixed Speed	Do Not Use	Do Not Use	Use
Variable + Bulk (00)	Use (top priority)	Do Not Use	Use
Variable + Low Latency (01)	Do Not Use	Use (top priority)	Use
Variable + Mixed (10)	Use (top priority)	Use (top priority)	Use
Variable + Intermediate (11)	Use (top priority)	Use	Use

Table C-25. Rules for Interpreting SENDER-APPROACH EOW

C.7.6.3.2. **Operation With Edition 3 Peers**

To communicate receiver recommendations to a peer that supports Edition 3, the ED3-BASIC-RATE (Type 1) EOW defined in Section C.6.2 **shall** be used. The ED3-BASIC-RATE EOW **shall not** be sent to peers known to support Edition 4 (or subsequent).

There are two key recommendations that can be derived from this EOW:

1. Transmission speed, with options for standard HF data rates. This speed should be treated as a recommended speed for maximum throughput and as a recommendation to not use a faster speed.

2. Interleaver. This should be treated as a recommended interleaver to use

C.7.7. Service Mapping onto Modem & TRANSEC Services

The Data Transfer Sublayer uses services from the layers below. The layer immediately below is the MAC layer. The MAC layer provides functions to control access to the channel. The interface to Modem and TRANSEC is specified as part of DTS, as the DTS is the primary user of these services. The MAC layer relationship to the DTS varies with the mechanism:
1. For CSMA (Annex K), the MAC layer simply manages timers that control access.
2. For WTRP (Annex L), the MAC layer has protocol, but this protocol uses the DTS peer protocols for data transfer of M_PDUs.

This layer below DTS/MAC may be one of two types of entity:

1. A Modem for HF or other frequencies; or
2. A TRANSEC device/layer which provides data encryption/decryption of all data transferred. This TRANSEC capability may also be used to provide COMSEC protection of user data,

There will be a direct data interface to the layer below, and three options to provide it are described in the following sections.

C.7.7.1. Modem Control Interface

The DTS can operate at fixed speed without access to modem control.

In order to change transmission speed and interleaver as specified in Section C.7.6, the DTS needs to be able to control the modem. In order to be able to make recommendations on transmission parameters as specified in Section C.7.6.3, a node needs access to SNR and other information from the modem.

A control connection from DTS to Modem is needed in order to achieve this. Where TRANSEC is used, this control connection needs to use a Crypto Bypass in order to communicate directly with the modem.

Modem control interfaces are not specified in STANAG 5066.

C.7.7.2. ALE Unit Control Interface

ALE Units are often provided in conjunction with HF Modem products. Control of ALE Units comes from two parts of the STANAG 5066 stack:

1. The Channel Access Sublayer controls use of 1:1 ALE. This access can be considered to pass down through the DTS.
2. MAC layer controls use of multi-node ALE.

A control connection from STANAG 5066 Stack to ALE Unit is needed in order to support these functions. Where TRANSEC is used, this control connection needs to use a Crypto Bypass in order to communicate directly with the modem.

ALE Unit control interfaces are not specified in STANAG 5066.

C.7.7.3. Data Interfaces

C.7.7.3.1. **Data Direct to Modem**

STANAG 5066 DTS can connect directly to a modem.

STANAG 5066 does not specify or require an interface here. MIL-STD-188-110D Appendix A specifies a protocol for data communication with an HF modem over TCP.   Use of this protocol in conjunction with STANAG 5066 to interface to a modem is **recommended**.

C.7.7.3.2. **Data to TRANSEC or Modem using Annex D**

STANAG 5066 Annex D ("Interface between Data Transfer Sublayer and Communications Equipment") defines an synchronous serial interface that can be used to connect directly to a modem or to a crypto device providing TRANSEC.

When using this interface for the DTS, the follow considerations are noted:

1. Because there is no clear start/stop timing for this interface, the DTS will generally need to pad the start and end of transmissions, in order to avoid data loss.

2. Because synchronous serial interfaces are bit-aligned, a receiving DTS **shall** treat an inbound data stream as a bit-stream and make use of the Maury-Styles D_PDU header to determine the byte alignment required by the DTS protocols

C.7.7.3.3. **Data to TRANSEC using Annex T**

STANAG 5066 Annex T ("TRANSEC Crypto Layer using AES and other Protocols") defines a TRANSEC layer with a peer communication protocol.   The DTS can use this directly, without need for padding.   The protocol specified is byte-aligned.

STANAG 5066 does not standardize an interface to the TRANSEC layer specified in Annex T.

C.8.   Extended D_PDU Summary

Annex L (HIGH-FREQUENCY WIRELESS-TOKEN-RING-PROTOCOL
(WTRP) REQUIREMENTS) has assigned Extended D_PDUs with values in range 1-8.   These are all used
for control purposes.   The state machines in this Annex have been written to support this usage.

Extended D_PDU Types 0, and 7-255 are reserved for use in future editions of STANAG 5066.   Note that new
assignments **may** require changes to the state tables of this annex.

C.9. Changes in Edition 4

This section describes changes in edition 4 Annex C, relative to edition 3 of STANAG 5066.

1. Operation over duplex links is made clear. A model is introduced, distinguishing single channel, duplex and broadcast operation. Sections are added to clarify operation.

2. Duplex used only for duplex operation. Edition 3 used the term Duplex to describe data flowing in both directions, which could be over duplex or half duplex link. This confusing usage removed.

3. Data rate adaption is only supported for auto-baud waveforms such as STANAG 4539 or STANAG 5069. Fixed speed operation over non-auto-baud waveforms is supported. Rationale for dropping the Edition 3 Data Rate Change mechanism that supports non-autobaud waveforms:

   o The Edition 3 mechanisms for data rate change are complex and fragile.

   o Data rate change with non-autobaud waveforms is not needed in modern deployments.

   o It is problematic to specify duplex operation of these mechanisms or use of these mechanism with ALE, both of which are key elements of Edition 4.

   o Many Edition 3 implementations do not support this option.

4. A number of new capabilities are introduced. These are added so that they will only be used in communication with peers that are known to support edition 4 or a subsequent edition. Mechanisms are defined to determine this support. The new mechanisms **shall not** be used with peers that are known to not support edition 4 or where peer capability is unknown. This ensures robust interoperability with older systems.

5. A padding D_PDU is introduced, which can increase resilience and performance by adding EOTs and EOWs to a transmission in space that could not otherwise be utilized.

6. A data rate selection process has been added, more appropriate to auto-baud waveforms. Interoperation with edition 3 is specified, as well as edition 4-specific capabilities that include support for WBHF speed recommendations/

7. Increased flexibility on D_PDU size choice which will improve performance, resilience and rapid handling of higher priority data.

8. Use of 16 bit frame sequence number, which is important to achieve good ARQ performance at WBHF and faster narrowband speeds.

9. An extension D_PDU that allows additional D_PDU types to be defined. This is used in Edition 4 Annex L and enables other new capabilities to be added.

10. Explicit prevention of in-order delivery for non-ARQ, as this can lead to system lock-up when data is lost.

11. Incorporation of implementation notes for STANAG 5066 Edition 3 Annex H (Implementation Guidance and Notes) at relevant points, with updates and extended notes based on more recent

experience.

12. Folding in technical elements (EOW additional values) and implementation notes from STANAG 5066 Edition 3 Annex G (Use of Waveforms at Data Rates Above 2400 bps).

13. New section clarifying D_PDU transmission rules.

14. New section clarifying mappings onto Modem and TRANSEC.

15. Clarification of use of ALE with reference to Edition 4 changes in Annex B and Annex J.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX D
STANAG 5066
Edition 4

Annex D: Interface between Data Transfer
Sublayer and Communications Equipment
(mandatory)
DRAFT FOR REVIEW: Version 1.0 of April 20th 2020

The interface between the Data Transfer Sublayer and the communications equipment **shall** [(1)] be as defined in this Annex or as specified in Annex T (STANAG 5066 TRANSEC Crypto Sublayer using AES and other Protocols). This annex **may** be used to communicate from Data-Transfer-Sublayer to a modem or to a cryptographic device providing TRANSEC services.

To accommodate requirements imposed by the communications equipment, an arbitrary number of 'pre-fill characters **may** [(1)] be transmitted prior to the first valid DPDU in a transmission interval. Likewise, an arbitrary number of 'post-fill' characters **may** [(2)] be transmitted following the last valid DPDU in a transmission interval. The value of the EOT assigned to each DPDU in the transmission interval **must** be computed accounting for any post-fill characters in the transmission interval. Extraneous characters **shall not** be inserted between any valid DPDUs in the transmission interval.

Note that this interface defines a bit stream, and that a receiving implementation will need to analyse the stream looking for DPDU Maury-Styles headers to obtain byte alignment.

The interface specifies a synchronous serial digital data interface.

The line-drivers and receivers for the interface **shall** [(4)] be configurable for either balanced or unbalanced connection, in accordance with EIA-232D/423 for unbalanced connections and EIA-422 for balanced connections.

With respect to functional roles on the interface, the Data Transfer Sublayer **shall** [(5)] be hosted in a Data Terminal Equipment (DTE).

The clock source for the data output from the DTE (i.e, DTE data out) on the interface **shall** [(6)] be either configurable or from the DCE (i.e, either the cryptographic equipment or the modem).

The clock source for the data input to the DTE (i.e., DTE data input) **shall** [(7)] be from the DCE (i.e, either the cryptographic equipment or the modem).

The interface **shall** [(8)] provide full hardware-level handshaking for flow-control, in accordance with any standard recommendations.

Compatability with MIL-STD-188-114 polarity, levels, and slew rates additionally **may** be required for interoperability with existing (cryptographic) equipment.

This Annex is broadly unchanged since Edition 1. Edition 4 has minor clarifications and alignment to the new Annex T.

Annex F: SAP Assignment
(Mandatory)
DRAFT 1.0 of 14th April 2020

Annex F defines SAPs that are assigned to applications using STANAG 5066 and default priorities for these applications.

F.1. SAP Assignment

Every application running over STANAG 5066 needs to use a SAP (range 0-15) and SAPs cannot be shared. Table F-1 specifies SAP assignment, with references to STANAG 5066 Annexes or to other open specifications.

Application	SAP ID	Reference
Subnet Management	0	Annex S (SIS Access Protocol)
Character-Oriented Serial Stream (COSS)	1	Annex P (ACP 127 & Character-Oriented Serial Stream)
ACP 142 (Military Messaging)	2	Annex Q (ACP 142)
Reserved	3	
Reserved	4	
HF Operator Chat	5	Annex O (HF Operator Chat)
XMPP	6	XEP-0365 "Server to Server communication over STANAG 5066 ARQ" https://xmpp.org/extensions/xep-0365.html
ACP 142 (Email)	7	Annex Q (ACP 142)
Reserved	8	
IP Client	9	Annex U (IP Client)
XMPP over SLEP	10	XEP-0361 "Zero Handshake Server to Server Protocol" https://xmpp.org/extensions/xep-0361.html over S5066-EP3 "SIS Layer Extension Protocol (SLEP)" https://www.isode.com/whitepapers/S5066-APP3.html
MULE over SLEP	11	RFC 8494 "Multicast Email (MULE) over Allied Communications Publication (ACP) 142" https://tools.ietf.org/html/rfc8494 over S5066-EP3 "SIS Layer Extension Protocol (SLEP)" https://www.isode.com/whitepapers/S5066-APP3.html
Compressed File Transfer Protocol	12	Annex V (Compressed File Transfer Protocol)
Unallocated	13	
Unallocated	14	
HF Discovery, Ping and Traffic Load	15	S5066-APP2 "HF Discovery, Ping and Traffic Load" https://www.isode.com/whitepapers/S5066-APP2.html

Table F-1: SAP Assignment

When protocols listed in Table F-1 are used, the assigned SAP ID **shall** be used. Reserved SAP IDs were used in Edition 3. They may be allocated or move to Unallocated in future editions of this annex.

ACP 142 is assigned two SAPs, as it is generally preferable to operate Email and Formal Military Messaging on different SAPs.

When assigning additional applications to SAPs for a deployment, preference **shall** be given to Unallocated SAPs. Where this is not possible, assignment of another SAP not used in the deployment **may** be made.

F.2. Default Priorities

It is desirable that applications operate at consistent priorities for a deployment. For this reason, default priorities for the applications listed in Table F-1 are given in Table F-2. Deployments **may** choose to follow this recommendation or **may** make deployment-specific assignments.

Application	Default Priority
Character-Oriented Serial Stream (COSS)	5
ACP 142 (Military Messaging)	Based on Message Priority • OVERRIDE: 13 • FLASH: 11 • IMMEDIATE: 9 • PRIORITY: 7 • ROUTINE: 5 • DEFERRED: 3
HF Operator Chat	12
XMPP	10
ACP 142 (Email)	Based on Message Priority: • URGENT: 6 • Normal: 4 • Non-Urgent: 2
IP Client	6
XMPP over SLEP	10
MULE over SLEP	As for ACP 142. If MULE over SLEP is used for both Email and Military messaging, a second SAP should be selected.
Compressed File Transfer Protocol	4
HF Discovery, Ping and Traffic Load	1 A low value is chosen, to not interfere with operational traffic

Table F-2: Default Priority Assignments

Table F-2 sets out default priority assignments. ACP 142 (Military Messaging) has priorities derived from the military message priority. The defaults are assigned so that there is space above, below and in between each one. Other application assignments have been made relative to this.

F.3. Changes Since Edition 3

Most of Edition 3 Annex F has been moved to other Annex. Annex F retains its core function of specifying SAP assignment. The SIS Access Protocol, which was specified in Annexes A and F in Edition 3 is now specified in Annex S. Most of the other applications specified in Edition 3 Annex F have their own Annexes in Edition 4. This improves modularity and maintainability of STANAG 5066.

There are only 16 SAPs available, and it is anticipated that this will lead to requirements for careful allocation in some deployments. For this reason:

1. Edition 3 assignments that are believe to be un-used are removed.

2. Assignments **shall** be followed if an application is used, but **may** be re-assigned for other use.

The following protocols specified in Edition 3 do not have annexes in Edition 4, and do not have SAP assignments:

1. HMTP. Deployments have shifted to CFTP, which offers compression.
2. HFPOP. No implementations known, and future use unlikely.
3. Ether Client. Only experimental implementations known, and future use unlikely.

BASIC FILE TRANSFER PROTOCOL (BFTP) and File-Receipt Acknowledgement Protocol (FRAP) are not specified in Edition 4.

UDOP and RCOP are not assigned SAPs, as the lack of protocol to support SAP sharing means that this approach has not been taken operationally.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

ANNEX J - GENERAL REQUIREMENTS FOR ENHANCED MEDIA-ACCESS-CONTROL (MAC) CAPABILITIES (INFORMATIVE)
DRAFT FOR REVIEW: Version 1.0 of March 31st 2020

J.1 INTRODUCTION

The channel-access control capability of STANAG 5066 Annex B Edition 1 provides mechanisms (i.e., the CAS-1 linking protocol) to establish a point-to-point link (or links) for data communication. The CAS-1 protocol belongs to the class of link request/accept protocols that are effective at resolving the hidden-terminal problem in wireless point-to-point scenarios, by ensuring peer communication only.

This annex introduces modes for enhanced media-access control capability for HF data communication networks, and the prescribed method in which they are used with other STANAG 5066 capabilities. These channel-access modes extend or modify, but do not replace, the channel-access and link-control mechanisms defined in Annex B of this STANAG.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

J.2   MEDIA-ACCESS-CONTROL SUBLAYER (MACS)

Enhanced Media-Access Control capabilities are defined in the context of an augmented model of the HF Subnetwork's protocol stack shown in Figure J-1.
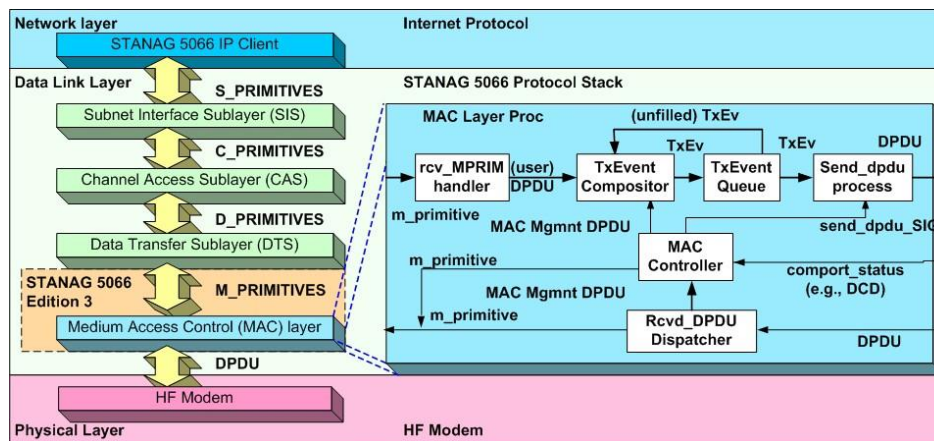
Figure J-1 — Augmented Model of the HF Subnetwork Protocol-Stack

The augmented model adds a Media-Access Control Sublayer (MACS) and inserts it below the Data-Transfer Sublayer of Annex C. Additional functionality implementing enhanced media-access control may be contained in the MACS, as shown in the   figure.

The added media-access-control functionality:

- **shall** be based on the D_PDU message types defined in STANAG 5066 Annex C, <u>and</u>

- **may** use in addition new D_PDU types implemented for media-access control that **shall** conform to the message-definition rules and the Generic D_PDU Frame Structure and D_PDU field-element requirements of S'5066 Annex C Sections C.3.1 and C.3.2. New functionality   to implement enhanced media-access-control **shall** be confined to the D_PDU Type-Specific Header element and, if present, D_PDU Payload element of any new D_PDU   type.

To minimize impact on pre-existing functionality defined in Annexes A, B, and C   for other layers of the HF subnetwork, new D_PDU types defined for media-access control **shall** only support peer-to-peer MACS-layer communication between   nodes.

Node-to-node communication between HF subnet clients and   peer-layer-to-peer-   layer communication supporting the functionality of Annexes A, B, and C **shall** continue to use   the D_PDU

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

types defined in Annex C. The S_PDU and C_PDU specifications of Annex A and   Annex B **shall** be unchanged by the MACs   functionality.
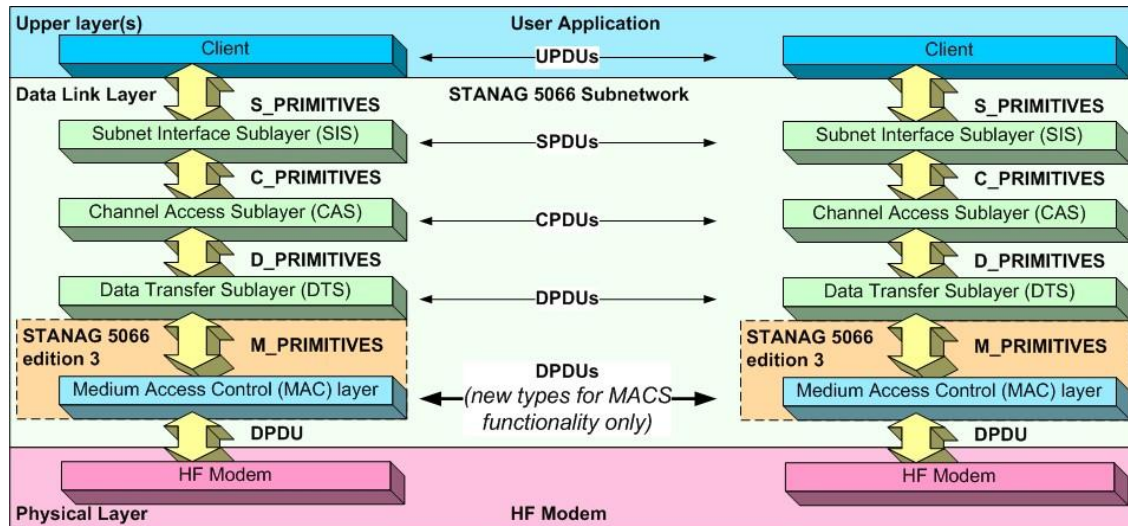
Figure J-2 — Peer-to-Peer Communication for Enhanced Media-Access   Control

MACS functionality may be tailored as outlined in section J.3, which defines options to implement a range of functionality for multi-node networks. Detailed functional and performance specifications from other Annexes are cross-referenced where appropriate.

J.3   SUMMARY OF MEDIA-ACCESS-CONTROL MODES

Four Media-Access Control Modes are defined:

1.   Point-to-Point Mode (P2P)

2.   Carrier-Sense Multiple Access (CSMA) Media Access

3.   Wireless Token-Ring Protocol (WTRP) Media Access

4.   Adaptive Time-Division Media Access (ATDMA)

These modes are summarized below. Implementation and performance requirements for these are standardized in the cross-referenced annexes.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

J.3.1 Point-to-Point (P2P) Media-Access Mode

The point-to-point (P2P) media-access mode enables two nodes to reserve the channel for their use in point-to-point communication. The default point-to-point channel- access control **shall** be implemented in accordance with STANAG 5066 Annex B. It **shall** also use CSMA following STANAG 5066 Annex K, noting that for two nodes only some simplification of that Annex K procedure is possible.

J.3.2 Carrier-Sense Multiple Access (CSMA) Media-Access Mode

The Carrier-Sense Media Access (CSMA) mode enables a set of nodes to share the channel in a multi-node network. CSMA is a form of media-access control based on a node's ability to listen to the channel (i.e., radio-frequency carrier) and use its local knowledge that the channel is clear to control its transmissions to avoid interference — also called collisions — with the transmissions of other nodes.

If a CSMA mode is implemented, it **shall** be implemented in accordance with STANAG 5066 Annex K.

The CSMA mode specified in Annex K conforms to the requirements of this Annex: it is implemented using only the message catalogue defined by STANAG 5066 Edition 1, and the basic capabilities of Annexes A, B, and C, with augmented requirements for Annex D to implement a Listen-Before-Transmit (LBT) mechanism with collision avoidance.

J.3.3 Wireless Token-Ring Protocol (WTRP) Media-Access Mode

The Wireless Token-Ring Protocol (WTRP) media-access mode enables a set of nodes to share the channel in a multi-node network. Token-ring or token-bus multiple access is a form of media-access control based on ownership of a 'token' that grants the right to transmit (RTT) on the channel. A token holder transmits until it no longer has data to send, or until its right-to- transmit timer expires, and then it passes the RTT token to its successor. Adaptation of each node's allocated channel capacity to the offered traffic load occurs automatically with the WTR mode.

The Wireless Token-Ring Media-Access mode **shall** be implemented in accordance with STANAG 5066 Annex L. The Wireless Token-Ring Protocol (WTRP) defined in that Annex is in two parts:

- a D_PDU message design for the management tokens exchanged by network nodes, and

- the algorithms used to create, maintain, and repair the ring (i.e., transmission sequence) of nodes in the network.

WTRP's D_PDU message design conforms to the requirements of this Annex. It uses new D_PDU message types based on EXTENSION D_PDU that conforms fully to the requirements of Annex C, whose D_PDU- specific part provides the requisite data fields to satisfy the information exchange requirements of WTRP.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

J.3.4   **Adaptive Time-Division Media-Access (ATDMA) Mode**

The Time-Division Media-Access (TDMA) mode enables a set of nodes to share the channel in a multi-node network. TDMA divides the channel access into timeslots that are allocated to nodes in the network (or, even more generally, allocated to network services).

Fixed TDMA protocols are known to be inefficient in channel utilization and service times when the traffic offered by the timeslot owner (i.e., the node or service to which the time slot is allocated) is mismatched to the channel capacity provided by the time slot. This is particularly true when the node or service has insufficient offered traffic to fill the timeslot(s) it has been allocated.

Adaptive TDMA (ATDMA) permits variations in the timeslot allocation or length that allow the network to adapt to variations in the traffic offered by nodes and in their service requirements. To use ATDMA's adaptivity with STANAG 5066 would require new D_PDU management message types with which nodes coordinate timeslot usage and length information.

There is currently no STANAG 5066 Annex that specifies TDMA or ATDMA.

J.4   **COMPATIBILITY AND INTEROPERABILITY ISSUES**

Through their conformance with the generic D_PDU-message structure defined in Annex C the enhanced media access protocols defined and cross-referenced herein are compatible in the limited sense that their common message elements (e.g., Maury-Styles synchronization preamble, D_PDU type field, address fields, etc.) are recognizable regardless of which MACS mode is in operation or has been implemented by the node. Thus, all STANAG-5066-compliant implementations will be capable of decoding the field elements common to all the D_PDU messages. And, in particular, all compliant implementations will be capable of determining the source address of the node that sent the D_PDU and the type number of the D_PDU that was sent, even if the node does not implement the MACs protocol for which the D_PDU is used.

But the enhanced media access modes defined herein are not interoperable. There is no intent or expectation that a node implementing one of these enhanced media-access modes should be interoperable with a node implementing a different protocol. Rather, the intent and expectation is that standard operating procedures for network establishment will ensure that only nodes using the same media-access control protocol are network members. As this is a naïve view of what can happen in an operational network, this Annex defines provisions for nodes to discover the MAC-mode in use by other nodes, to, at the very least, recognize when they are using incompatible protocols and to take appropriate action to avoid mutual interference.

J.4.1   **MAC-Mode Discovery**

The processes by which nodes discover the media-access-control modes in use within a STANAG

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

5066 network are called MAC-Mode Discovery.

Nodes **should** implement MAC-Mode Discovery, which consists of the process elements defined here:

• Use of D_PDU Type 15 Warning Messages – the (mandatory) provisions of Annex C Section C.3.12 **shall** apply to the implementation of MAC-Mode Discovery, as further amplified below.

• Channel Analysis – nodes implementing a given media-access mode shall analyze the channel usage by other nodes to detect violations of the protocol that it implements, and take actions as noted further below. These actions include the use of Type 15 Warning messages to notify neighbouring nodes of the protocol incompatibility.

J.4.1.1   *Use of Type 15 Warning Messages*

A STANAG 5066-compliant node that receives a D_PDU message that is "unexpected or unknown" to it is required under the provisions of Annex C Section C.3.12 to send a Type 15 D_PDU (Warning) Message to the node that generated the unexpected or unknown D_PDU. This requirement is unchanged by the enhanced-media-access modes defined herein.

In particular, and with cross reference to Annex C Section C.3.12, a node that receives a D_PDU that is unexpected for the MACS protocol is it is using (this may be an indication that the sending node is using a different MACS protocol) or that unknown to it (also a potential indication of an incompatible MACS protocol) will determine the source address of the node that sent the unexpected or unknown D_PDU and send a Type 15 D_PDU (Warning) Message to that node:

• As required in Annex C, the Type 15 D_PDU (Warning) Message RECEIVED FRAME TYPE field **shall** indicate the Type number of the unexpected or unknown D_PDU.

• If the node does not recognize the received D_PDU, the Type 15 D_PDU (Warning) Message REASON WARNING SENT field **shall** be set to the value assigned to the reason "Unrecognised D_PDU type Received".

   o N.B.: This case applies to D_PDU types with subtypes, i.e., the Type 6 D_PDU (Management) that uses the Engineering Orderwire (EOW) Type field (as defined in Annex C) as a subtype indicator to distinguish variants of the Extended Management Message types: the receiving node may recognize the main type (i.e., Type 6 Management Message) but not the subtype (i.e., the particular Extended Management Message subtype designated by the EOW-type field type). In this case, the node **shall** treat the D_PDU as an unrecognized D_PDU, and send a Type 15 D_PDU (Warning) Message as above.

• If the node recognizes the received D_PDU, but it was unexpected for the protocol that it was using,

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

the Type 15 D_PDU (Warning) Message REASON WARNING SENT field **shall** be set to the value assigned to the reason "Invalid D_PDU Received for Current State".

- o N.B.: This case also applies to D_PDU types with subtypes as described in the preceding case. If the node receives a D_PDU of the proper type but if the subtype is unexpected, e.g., it designates an Extended Management Message defined for a MACS protocol that the node recognizes but is not currently executing, then in this case also, the node **shall** treat the D_PDU as an unexpected D_PDU, and send a Type 15 D_PDU (Warning) Message as above.

- In particular, a node executing the P2P or CSMA mode that receives any of the new D_PDU types defined in Annex L (for WTRP) (whether or not they are recognized to the receiving node) **shall** send a Type 15 D_PDU as specified above.

Use of Type 15 D_PDU (Warning) Messages of course is possible only when the node is configured in a compatible transmission mode; nodes configured for transmit-only operation (e.g., to providing exclusive support to a Broadcast Data Exchange Session as defined in Annex A) will not receive any D_PDUs that could trigger the warning condition, and nodes configured for receive-only operation (e.g., to receive the Broadcast Data) would not be capable of sending the warning messages and therefore disable the protocol-error-detection logic and warning-message- generation functionality.

J.4.1.2   *Channel Analysis*

Transmissions by nodes executing one MAC-mode will likely be uncoordinated with the transmissions of nodes executing another (i.e., they don't participate in the protocol), and mutual interference may occur.

Nodes executing a given MAC-mode **shall** analyze the D_PDU messages received from other nodes to detect violations of the protocol it is executing. Message analysis includes the following:

- processing of received Type 15 D_PDU (Warning) Messages in accordance with Annex C and as amplified by Section J.2.1.1.

- analysis and identification of the MAC mode in use by other nodes:

  - o transmissions that contain none of the new D_PDU types (recognizable to the receiving node) defined in Annex L (for WTRP) **should** be assumed to originate from a node that is currently operating in P2P or CSMA modes, as these modes (i.e., P2P and CSMA) do not use any new D_PDUs for their operation. This indication is one of the current operating mode, and not of capability.

  - o Transmissions that contain Type 15 D_PDU (Warning) Messages that warn ignorance of any of the new D_PDU types defined in Annex L (for WTRP), i.e., that have a RECEIVED

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

FRAME TYPE field defined for one of the new D_PDU types and that denotes a REASON WARNING SENT field as unrecognized, **should** <u>be assumed to originate from a node that is capable of operating only in P2P or CSMA modes</u>. The presumption **should** be that node is a legacy (STANAG 5066 Edition 1 compliant) implementation incapable of executing either WTRMA or ATDMA.

o Transmissions that contain Type 15 D_PDU (Warning) Messages that warn of unexpected use of any of the new D_PDU types defined in Annex L (for WTRP), i.e., that have a RECEIVED FRAME TYPE field defined for one of the new D_PDU types and that denotes a REASON WARNING SENT field as unrecognized, **should** be assumed to originate from a node that is operating in a different MACS mode than the node to which the Type 15 D_PDU (Warning) Message is addressed.

J.4.2   Embedding Broadcast / Multicast Traffic

Receive-only nodes (e.g., nodes in an emission-control [EMCOM] status) make no transmissions that can interfere with the operation of multi-node network, whatever MACS mode it uses. This may be exploited by nodes that do transmit in a multi-node network to embed broadcast or multicast traffic, a capability that is fully consistent with the intent of Annex A Section A.1.1 (from Edition 2 and later of this STANAG).

J.5   Mappings onto Channel and use of ALE

The MAC layer **may** be mapped onto the underlying HF channel in one of the three ways set out below.

J.5.1   Fixed Frequency

In this mode, a fixed frequency agreed by all nodes is used. If STANAG 5069 Wideband HF is used, a fixed bandwidth is also chosen. This **shall** be the same for all nodes on the network. This is a simple and robust choice that is effective for some deployments, such as a network using HF surface wave.

J.5.2   Scheduled Frequency

A variant on fixed frequency is to have an agreed schedule of frequency use that is followed by all nodes. This can allow for adaptation to conditions in a simple manner, for example to use different daytime and nighttime frequencies.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX J
STANAG 5066
Edition 4

J.5.3   ALE for All Nodes

STANAG 5066 Annex B defines channel access mechanisms that set up ALE links on demand for point to point and multicast links.

An alternate approach is to use ALE for all nodes on the channel, which is managed at MAC layer. By maintaining an open channel, it enables all nodes on the channel to communicate with CSMA or WTRP. It will prevent the situation where a pair of nodes are connected on a channel, which blocks communication with other nodes.

Use of ALE at MAC level means that a channel can be kept open for all nodes and that the best channel at a given time can be chosen. It is important that the ALE channel is closed and re-opened from time to time, in order to ensure choice of best channel and to bring on nodes that were not available on the previous ALE setup.

J.6   Changes in Edition 4

1.   Changed to include networks of any number of nodes. Edition 3 title and introduction suggest multiple nodes, but this is contradicted by much of the text which covers two node networks.

2.   Removed normative references to Annex M, which is a placeholder.

3.   Require use of CSMA for point to point. Edition 3 notes that MACS adds no functionality which is contradicted by footnote explaining why MACS should be used.

4.   Added section on mapping to channel and use of ALE.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

ANNEX K – HIGH-FREQUENCY CARRIER-SENSE MULTIPLE-ACCESS PROTOCOL (OPTIONAL) DRAFT FOR REVIEW: Version 1.1 of April 2nd 2020

K.1   INTRODUCTION

This Annex specifies a High-Frequency Carrier-Sense Multiple Access (CSMA) with Collision Avoidance protocol [1] for STANAG 5066 in multi-node single-frequency networks.

The HF CSMA protocol introduces no new DPDU types to the STANAG 5066 catalogue. There is no explicit peer-to-peer communication required by the CSMA protocol. CSMA media access control relies on a node's local information of HF channel activity and inferences of the activity ― or lack thereof ― of other nodes.

In order to be effective, the options used in this annex and the values for timers need to be set consistently for all nodes on the network.

Annex K is Optional. However if CSMA or any other form of Listen Before Transmit is implemented with STANAG 5066, it **shall** conform to Annex K.

Section K.2 of this Annex presents an overview of the protocol and provides a definition of terms. Details of the protocol, its state diagram, and parameter values are specified in section K.3

K.2   OVERVIEW: CARRIER-SENSE MULIPLE-ACCESS PROTOCOLS

Definitions and management concepts are introduced below prior to detailed specification of the protocol in later sections.

K.2.1   Background

[1]The protocol specified in Edition 3 was based on a presentation to the High-Frequency Industry Association in 2002 by Robert McFarland of Rockwell Collins., "Collision Avoidance using STANAG 5066 in a Network Environment", with additional input from Michael Stringer of Harris Corporation. This used a Jitter mechanism so that many nodes on a network could reduce the chance of collision. Standardizing this enables multi-vendor deployments.

The update for Edition 4 uses work from Isode specified in the paper "Slotted Option for STANAG 5066 Annex K". This introduced the following improvements:

1. An optional "slotted" mechanism where each node is assigned a slot. This gives faster switching and higher resilience for small networks.

    a. The "jitter" mechanism remains an option for larger networks, although modern large networks will generally use ALE rather than fixed frequency.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

2. Optimized switching when just two nodes are communicating. This removes any residual benefits of operating without CSMA (Annex K) or WTRP (Annex L).

3. Use of different timings when STANAG 5066 EOT is detected to enable faster switching than is possible with a single timer.

4. Optimizing repeat transmissions from a single node, which is important for non-ARQ bulk protocols such as ACP 142.

This combination leads to Annex K providing an alternative to WTRP (Annex L) which will generally provide better performance for a lightly loaded network.

K.2.2   Definitions

The following terms are used in the specification of the High-Frequency Carrier-Sense Multiple Access (CSMA) protocol.

K.2.2.1   Stations and Nodes

The terms "*station*" and "*node*" are used interchangeably to describe the communication entities on the shared HF channel.

K.2.2.2   Collisions

Channel collisions ― or, simply, *collisions* ― are simultaneous or overlapping transmissions by two or more nodes that interfere with each other, preventing reception by another node.

K.2.2.3   Listen-Before-Transmit (LBT)

Listen-before-transmit is the action a node takes to ensure that the channel is unoccupied and free of activity before it attempts to transmit.

K.2.2.4   DCD

DCD is the *Data-Carrier-Detect* signal provided by the communications equipment interface. A node senses that an HF radio carrier is present when DCD is *true* (or ― using the nomenclature of some interfaces ― asserted). Communications equipment that senses an idle channel will provide set DCD = *false* (noted in this annex as the !DCD signal). DCD does not always give a clear indication of when a transmission ends, due to channel fades. Note that DCD requires an interface from modem to STANAG 5066, which might need a crypto bypass.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

K.2.2.5   VDCD

VDCD is a *Virtual Data-Carrier-Detect* signal, derived from reception of valid DPDUs. Observation of DPDU header's End-of-Transmission (EOT) field (defined in Annex C Section C.3.2.3) **shall** be used to extrapolate into the future and predict the time at which a channel will be idle. The EOT mechanism provides a robust mechanism to determine end of transmission. It is preferred to DCD, and when VDCD is available it enables faster switching.

K.2.2.6   Contention Interval

The *contention interval* is the period of time during which nodes attempt to access the shared HF channel.

K.2.2.7   Collision Avoidance

Collision Avoidance is a strategy that nodes use during the contention interval to increase the probability that some contending node successfully accesses the channel.

K.2.2.8   Node States

A node may be in one of the following states as it executes the CSMA protocol:

- *Offline* state (OFFLINE) — the *offline* state is a state in which a station acts as if it were physically offline, i.e., it can neither transmit nor receive;

- *Sensing* State (SENSE) — the *sensing* state is a state in which a station monitors channel activity (using the DCD or VDCD signals), waiting for it to become idle.

- *Listen-before-Transmit-Wait* State (LBT_WAIT) — the *Listen-before-Transmit-Wait* state is a state in which a station waits to determine if another node has transmitted on the idle channel.

- *Contention-Wait* State (CONT_WAIT) — the *contention-wait* state is a state in which the station waits a random time before it may transmit. Nodes wait a random time during this state while continuing to sense channel activity to avoid collisions.

- *Linking* State (LINKING) — the *Linking* state is a state in which a station can transmit data.

K.2.2.9   Timers

The following timers control operation of the CSMA protocol:

- LBT_WAIT_TIMER – the LBT_WAIT_TIMER controls the maximum length of time a node will wait on an idle channel before transitioning to the contention-wait state. This value is determined from a number of parameters.

- CONT_WAIT_TIMER – the CONT_WAIT_TIMER controls the length of time a node (i.e., a contending node) waits to access the channel before it starts to transmit.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

K.2.2.10   Scalar Control Parameters

The following scalar parameters control operation of the CSMA protocol:

- LBT_WAIT_TIMER_VALUE – the waiting time set for the LBT_WAIT_TIMER determined from a number of parameters;

- LBT_WAIT_TIMER_VALUE_DCD – the waiting time set for the LBT_WAIT_TIMER when end of transmission is determined by DCD;

- LBT_WAIT_TIMER_VALUE_VDCD – the waiting time set for the LBT_WAIT_TIMER_GC when end of transmission is determined by VDCD;

- CONT_WAIT_TIMER_VALUE – the waiting time set for the CONT_WAIT_TIMER; this is a computed value that depends on other parameters and whether Jitter or Slotted is used.;

- NUM_CONT_SLOTS – the number of slots defined. For slotted, this **shall** be greater than or equal to the number of nodes on the network;

- CONT_SLOT_WIDTH – the duration of each contention slot.

- NODE_SLOT_POSITION: For slotted, a per node configuration in the range 1 to NUM_CONT_SLOTS. Each node **shall** have a different value configured.

K.2.3   Concept of Operations

A likely occurrence in multi-node HF subnetworks is that nodes need to communicate at the same time and may attempt to do so. Uncontrolled attempts to access and transmit on the same channel can lead to channel collisions when these transmissions occur at the same time. Mitigating the effects of collisions requires retransmissions and delays that lower channel throughput, but, when retransmissions remain uncontrolled, can further decrease throughput. Repeated collisions on an uncontrolled channel can degrade the network severely. Carrier-Sense Multiple Access (CSMA) concepts that include collision-avoidance (CA) can provide simple yet effective mechanisms for a (small) number of nodes to share a single- frequency HF subnetwork when they each have traffic to send.

The concept of operations for the CSMA protocol follows.

From an offline, non-operating state a node starts the protocol by entering a carrier- sensing state, where it listens on the channel for a radio-frequency carrier. The mechanisms a node uses for carrier-sense may vary but this Annex assumes that the carrier is sensed using either the Data-Carrier-Detect (DCD) signal available from the communications equipment, or from a Virtual Data-Carrier-Detect signal generated by tracking the EOT-field of received DPDUs contained in previous transmissions.

A node with no data queued for transmission remains in the carrier-sensing state.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

A node with queued data to send that senses an idle channel enters the Listen- Before-Transmit state, and sets a timer — the LBT_WAIT_TIMER — that controls its exit from the state. The timer is restarted whenever the carrier is detected. If the LBT_WAIT_TIMER expires and the node still has queued data, the node enters a contention state. At any time in the Listen-Before-Transmit state, if the node no longer has any queued data, the node returns to the carrier-sensing state. Note that the LBT_WAIT_TIMER timer value specifies the maximum time a node will wait on an <u>idle</u> channel. The node will wait for a longer time — potentially a much longer time — on a busy channel because the LBT_WAIT_TIMER is restarted whenever carrier is sensed.

Nodes wait before contending for the channel and transmitting because of the time delays associated with the carrier-sense mechanism. A node waits to determine that the channel truly is idle. Detection of a busy channel — i.e., declaration of the DCD or VDCD signals — is not instantaneous. A node's transmissions will have a propagation delay between itself and any receiver. Then, DCD is declared by a receive modem on declaration of valid data and it takes the modem's decoder some number of interleaver frames to make this declaration, introducing additional delay. Additionally, for secure systems that use crypto between the HF modem and the STANAG 5066 system, the modem DCD either needs to use a crypto bypass or to be relayed through the cryptographic equipment, which may introduces additional delay. These delays influence the length of the listen-before-transmit interval that enables nodes to avoid — but not completely prevent — collisions on the channel.

On completion of the Listen-before-Transmit interval i.e., when the LBT_WAIT_TIMER expires, a node with queued data to transmit enters a contention state, where it uses one either a jitter or a slotted mechanisms set out below.

For both mechanisms, the contention-state's random waiting interval is quantized. It consists of a number of slots, each of a fixed width. The slot width is specified to insure that all nodes in the network can detect a node's transmission before the beginning of the next slot.

The slot width **shall** be chosen so that it is greater than the channel-sensing delays — i.e., the sum of the modem-, crypto-, propagation- and other delays.

K.2.4 Jitter Mechanism

With the jitter mechanism, a node selects a random interval to wait before it enters the transmit state. A node picks a random time to wait by picking a random slot to start its transmission.

The random waiting period during the contention interval is intended to avoid collisions — the carrier-sense and listen-before-transmit states tend to synchronize the access attempts of multiple nodes and the random waiting time during the contention interval forces, with some probability of success, the attempts to occur at different times. With good choice of slot width, the probability of collision can be reduced to the probability that two or more contending nodes pick the same contention slot.

Quantizing the contention-state's waiting time and constraining nodes to attempt transmissions only at slot boundaries will reduce average wait time by reducing collisions. Unconstrained transmission- attempt times during the contention interval increases the possibility of collisions between contending nodes.

K.2.5 Slotted Mechanism

With the slotted mechanism, there is a slot assigned for each node on the network, and the total number of

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

slots is equal to the number of nodes on the network.

Advantages of the slotted mechanism are:

1. Each node will transmit in its own slot, so there is no risk of nodes picking the same slot.

2. For small networks it is safe to use a small number of slots.

   a. Note that for a two-node network, slots are not needed at all, because of a mechanism described later:

Disadvantages of the slotted mechanism are:

1. It requires more coordinated configuration than is needed by the jitter mechanism.

2. It is inefficient for larger networks, where Jitter is preferable.

3. It is not "fair" as priority is given to nodes with earlier slots. This may be an advantage in some deployments.

   a. The unfairness only becomes noticeable under high load. For a highly loaded channel, WTRP (Annex L) is likely to be preferable.

K.2.6   Optimizing for Two Active Nodes

Use of CSMA naturally introduces delays between transmissions in order to avoid collisions, while allowing multiple nodes to share a network. A common scenario on a lightly loaded network, for which Annex K is targeted, is that active communication is restricted to a pair of nodes. Annex K optimizes for this scenario, by allowing a node to transmit at the earliest safe point, when it knows that there is only a pair of nodes communicating. This gives useful performance optimization for this common scenario.

Note that either active node can introduce communication with a third node at any time, by establishing a CAS-1 link. However, this mechanism does block communication between other nodes until the pair of nodes has finished communicating. This is not expected to be a problem on target networks (lightly loaded).

K.3   SPECIFICATION AND PROTOCOL

The CSMA protocol is specified below in the following sections:

- K.3.1 Overall State Diagram , which presents the protocol as a state graph with directed transitions and transition events that may be conditional or unconditional;

- K.3.2 State Specifications, which presents detail of each state and outbound-transition tables that govern the actions that take place when the node is in the state, and the events that trigger transitions to the next state.

K.3.1   Overall State Diagram

Figure K-1 below shows the simplified CSMA state diagram in which only the states and transition events are shown. Each state is specified in detail in the following section.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4
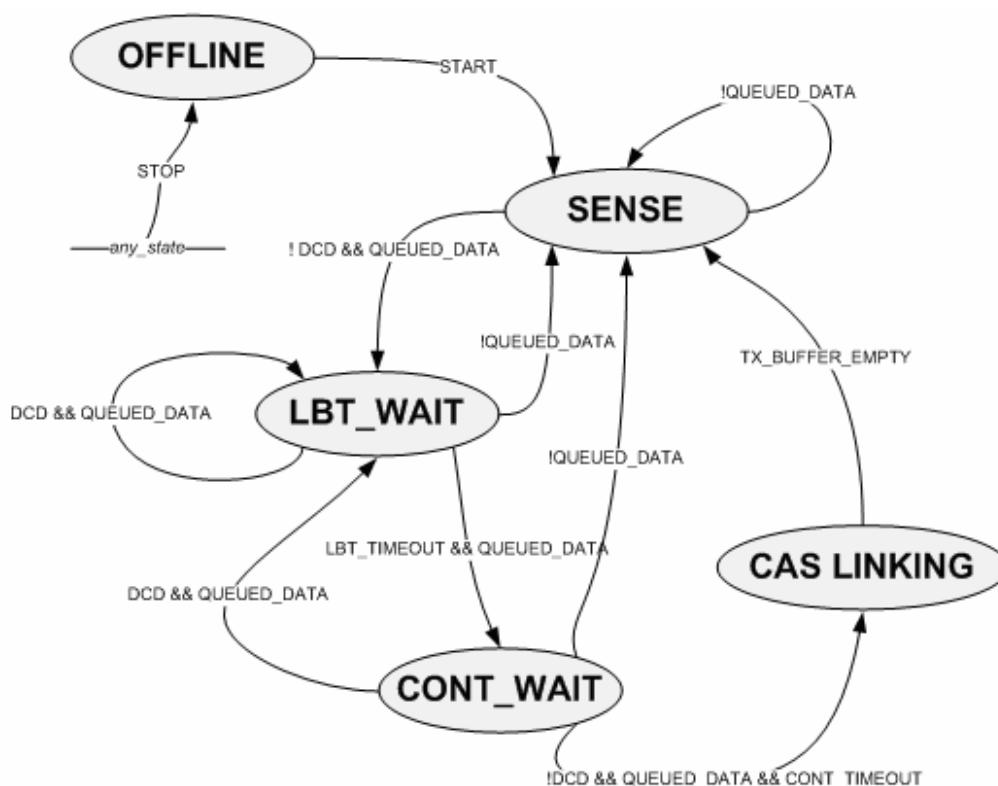
Figure K-1 — CSMA state diagram.

K.3.2   State Specifications

Specifications for CSMA operation for each state are provided below. For specification-purposes, transitions from each state to the next state shall be controlled by the accompanying outbound-transition tables, which specify:

- The current state,
- The event that triggers the transition,
- The action that shall be taken as a result of the transition,
- The next state to which the protocol transitions,
- The timer (if any) that is started.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

K.3.2.1   Offline State (OFFLINE)

A node in the OFFLINE state **shall** neither send nor receive data.

Outbound transitions from the OFFLINE state **shall** conform to the table below.

Table 1 — OFFLINE-State Outbound-Transition Table

state	event	Condition	Action	next state	start timer
OFFLINE	START	start event received from the subnetwork management function	Neither send nor receive	SENSE	*none*

A node **shall** transition to the SENSE state when it receives a START signal from the subnetwork management function (i.e.., when the nodes starts operating).

K.3.2.2   Carrier-Sensing State (SENSE)

A node in the SENSE state **shall** listen for an idle channel, which is declared whenever the DCD signal is false (i.e, a !DCD signal).

Outbound transitions from the SENSE state **shall** conform to the table below.

Table 2 — SENSE-State Outbound-Transition Table

state	event	Condition	action	next state	start timer
SENSE	!DCD (idle channel)	data QUEUED for transmission		LBT_WAIT	LBT_WAIT_TIMER
SENSE		no data for transmission		SENSE	*none*

A node with no data queued for transmission **shall** remain in the SENSE state regardless of channel activity.

A node with data to send **shall** wait in the SENSE state until it detects an idle channel, and then **shall** transition to the LBT_WAIT state. The value that LBT_WAIT_TIMER is to set depends on a number of conditions.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

1. LBT_WAIT_TIMER **shall** be set to zero if the node determines that the only active CAS-1 link is between the node and one other node. This is determined if one the following conditions are true:

   a. If the received transmission contains D_PDUs directly addressed to the local node and there are no D_PDUs directly addressed to other nodes. Note that the received transmission **may** contain non-ARQ D_PDUs addressed to broadcast and/or multicast destinations; or

   b. A transmission is received, but the local node is not able to parse any D_PDUs in the transmission, and the local node is known to be a peer in all known CAS-1 transmissions. Under these conditions it is able to transmit safely, as no other node is expected to transmit at this point. In order to determine this, a node needs to monitor the status of all CAS-1 links on the channel (not just the CAS-1 links that the node is involved in).

   The first condition **shall** be checked.   The second condition **may** be checked.

2. LBT_WAIT_TIMER **shall** be set to LBT_WAIT_TIMER_VALUE_DCD if the !DCD event was determined by DCD and not by VDCD.

3. LBT_WAIT_TIMER **shall** be set to LBT_WATI_TIMER_VALUE_VDCD if !DCD event was determined by VDCD (EOT).

K.3.2.3   Listen-Before-Transmit-Wait State (LBT_WAIT)

A node in the LBT_WAIT state **shall** wait while continuing to sense channel activity, to ensure that the channel is truly idle.

Outbound transitions from the LBT_WAIT state **shall** conform to the table below.

Table 3 — LBT_WAIT-State Outbound-Transition Table

state	event	condition	action	Next state	start timer
LBT_WAIT	DCD (busy channel)	data QUEUED for transmit		LBT_WAIT	LBT_WAIT_TIMER (restarts timer)
LBT_WAIT	LBT_WAIT_TIMER timeout event	data QUEUED for transmission	compute CONT_WAIT_TIMER_VALUE	CONT_WAIT	CONT_WAIT_TIMER
LBT_WAIT		no data for transmission		SENSE	*none*

At any time while in the LBT_WAIT state, detection of a busy channel (i.e., receipt of a DCD or VDCD signal) **shall** restart the LBT_WAIT_TIMER with a value of LBT_WAIT_TIMER_VALUE_DCD; the node **shall** remain in the LBT_WAIT state in this case.

At any time while in the LBT_WAIT state, detection of an empty data queue (e.g., it may have been emptied

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

by a management function) **shall** force a transition to the SENSE state.

If the LBT_WAIT_TIMER expires, CONT_WAIT_TIMER is set and calculated in the following manner.

1. If the Jitter approach is taken, CONT_WAIT_TIMER **shall** be set to multiple of CONT_SLOT_WIDTH, multiplied by a random number in the range 0 to NUM_CONT_SLOTS-1.

2. If the slotted approach is taken and the last transmission was not made by the local node, CONT_WAIT_TIMER shall be set to:

   (NODE_SLOT_POSITION -1) * CONT_SLOT_WIDTH

3. If the last transmission was made by the local node, CONT_WAIT_TIMER shall be set to:

   NODE_SLOT_POSITION * CONT_SLOT_WIDTH

The third condition is to address sending Non-ARQ data. When sending ARQ data, it is expected that another node will transmit next. When sending Non-ARQ, this may not be the case.

K.3.2.4   Contention-Wait State (CONT_WAIT)

Outbound transitions from the CONT_WAIT state **shall** conform to the table below.

Table 4 — CONT_WAIT-State Outbound-Transition Table

state	event	condition	action	next state	start timer
CONT_WAIT	DCD (busy channel)	data QUEUED for transmission		LBT_WAIT	LBT_WAIT_TIMER (restarts timer)
CONT_WAIT	!DCD && CONT_WAIT_TIMER timeout event	data QUEUED for transmission		LINKING	*none*
CONT_WAIT		no data for transmission		SENSE	*none*

Detection of DCD (i.e., detection of a busy channel) while in the CONT_WAIT state is taken as an indication that the node has 'lost' the contention round to another, and thus, if the node still has data to transmit, it **shall** restart the LBT_WAIT_TIMER with value of LBT_WAIT_TIMER_DCD and transition to the LBT_WAIT state to wait once again for an idle channel.

Expiration of the CONT_WAIT_TIMER while the channel remains unoccupied (i.e. without any detection of DCD) is taken as an indication that the node has 'won' the contention round and thus, in this case and if the node still has queued data for transmission, the node **shall** transition to the LINKING state.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

At any time while in the CONT_WAIT state, detection of an empty data queue (e.g., it may have been emptied by a management function) **shall** force a transition to the SENSE state.

K.3.2.5   CAS Linking State (LINKING)

Outbound transitions from the LINKING state shall conform to the table below.

Table 5 — LINKING-State Outbound-Transition Table

state	event	condition	action	next state	start timer
LINKING	(TX_BUFFER_EMPTY) or (TX_TIME > MAX_TRANSMIT_INTERVAL)			SENSE	*none*

A node in the LINKING state transmit until its transmit-buffer is empty (i.e., TX_BUFFER_EMPTY == TRUE) or until the time it has transmitted exceeds the maximum transmit interval allowed to it (N.B.: a node's maximum transmit interval can be limited by various considerations, whether it is the 127.5 second limitation imposed by the EOT field size, or the time it takes to transmit the maximum number of unacknowledged DPDUs allowed by the ARQ protocol, or some other lesser time imposed by the subnetwork management function.

K.4   SETTING SCALAR CONTROL PARAMETERS

Operation of the CSMA protocol is controlled by the scalar parameters listed in the table below.

Table 7 —Default values for CSMA Scalar Parameters

Parameter Name	Default Value	Units	Default-Value Name; Comments
CONT_SLOT_WIDTH	3	Seconds	CONT_SLOT_WIDTH; optimization of this value requires that it be a function of the modem preamble duration, data-rate, interleaver duration.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX K
STANAG 5066
Edition 4

NUM_CONT_SLOTS	16	integer	For a slotted configuration, NUM_CONT_SLOTS should be set to the number of nodes on the network. For Jitter; the value selected is a balance between probability of one and only one node selecting the winning slot and acceptable delay (eg. 3 nodes, 90% probability, 16 slots).
LBT_WAIT_TIMER_VALUE_DCD	30	seconds	This needs to be set to a long value as the measurement may be premature due to fades
LBT_WAIT_TIMER_VALUE_VDCD	3	seconds	A short time is recommended to allow for potential errors in EOT calculation. This also needs to be long enough to give nodes waiting for this sufficient time to detect a node that responds immediately.

K.5   Changes since Edition 4

Edition 4 introduces a number of new timers and scalars, to meet the goals set out in Section K.2.1, which extend the capabilities of Edition 3.

Because Annex K requires that all nodes are configured consistently, this does not introduce interoperability issues. A deployment with a mix of Edition 3 and Edition 4 systems will need to be configured with Edition 3 settings only.

Status is changed from Informative to Optional. As Annex K controls timers it impacts interoperability and is not merely Informational. This change also brings it in line with Annex J.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

ANNEX L - HIGH-FREQUENCY WIRELESS-TOKEN-RING-PROTOCOL (WTRP) REQUIREMENTS
DRAFT FOR REVIEW: Version 1.1 of April 8th 2020

L.1.   INTRODUCTION

This Annex specifies a High-Frequency Wireless Token Ring Protocol (WTRP) for enhanced media access control within single-frequency multi-node radio networks using STANAG 5066. The protocol is in two parts: a message design for the management tokens exchanged by nodes in the radio network, and the algorithms used to create, maintain, and repair the radio-transmission sequence (i.e., the virtual ring) of nodes in the network.

There is a token that represents right to transmit, which is transferred using a STANAG 5066 EOW (Engineering Order Wire) message. This means that the token can be transmitted without any protocol overhead and can be repeated for resilience.   Additional messages to manage the token ring are sent using STANAG 5066 EXTENSION D_DPUs as specified in Annex C.

The High Frequency Wireless Token Ring Protocol (WTRP) is a self-organizing Medium Access Control (MAC) protocol for HF wireless networks. The MAC protocol by which mobile stations can share a broadcast channel is crucial in wireless networks, especially for HF wireless networks where bandwidth is considerably lower than other types of wireless networks. WTRP is a MAC protocol derived from Ergen et al [1] and tailored for low-speed wireless networks. WTRP provides bounded latency. WTRP is efficient in reducing the number of retransmissions due to collisions and is fair in that each station uses the channel at least once for each ring traversal.

All stations are connected on a single channel with common waveform operating on the same frequency. For waveforms with variable bandwidth a common bandwidth must be used by all stations.   The frequency and bandwidth may be fixed or may be negotiated by ALE for all stations. For autobaud waveforms, transmission speed and interleaver may vary during operation or fixed values may be configured.

WTRP requires that stations in a ring take turns to transmit for a specified maximum amount of Time, with an order of transmission that includes each node at least once. WTRP is robust against single node failure. WTRP is different from its parent protocol [1] in that it provides the notion of self-rings, that it supports connected networks with arbitrary connectivity, that there is no ring owner, and that stations can transmit more than once per ring cycle.

This Annex of STANAG 5066 is organized as follows.

- · Section L.2 gives an overview of WTRP;
- · section L.3 specifies the STANAG 5066 message design for WTRP management messages;
- · section L.4 gives the complete WTRP state diagram and description for each state in WTRP;
- · section L.5 specifies parameters and timer selection for ring-management and operation;
- · section L.6 specifies requirements for token and message transmission.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.2. OVERVIEW: WIRELESS TOKEN RING PROTOCOL

Token-Ring definitions and management concepts are introduced below prior to detailed specification of the protocol in later sections.   This protocol is based on work by Mustafa Ergen, Duke Lee et. al., "Wireless Token Ring Protocol", University of California, Berkeley, CA 94720, USA.

L.2.1. Definitions

The following terms are used in the specification of the High-Frequency Wireless Token Ring Protocol (WTRP).

L.2.1.1. Stations and Nodes

The terms "station" and "node" are used interchangeably to describe the communication entities on the shared transmission medium.

L.2.1.2. Token, Messages, Rings and Ring Members

The WTRP protocol is a Medium Access Control (MAC) protocol. The task of this protocol is to schedule the access of two or more stations connected to the same physical medium, nominally an HF wireless network. Messages exchanged between stations for WTRP control are called *messages*.

The WTRP protocol organizes stations in such a manner that they rotate a *right-to-transmit Token (*or just *token)* among stations connected to the same physical medium. Only a *station* that receives the *right-to-transmit token* has the right to transmit user data. This *station* is then the *token-holder*. In normal operation there should only be one *token-holder*.

A set of stations sharing the same *right-to-transmit token* is defined as a *ring*, or *virtual ring.* A station participating in a ring is called a *ring member*.

When a station starts initially it is not a member of any ring. It will only be able to become a *ring member* if there is at least already one station connected to the same physical medium. If the station connected to the medium finds out there is no existing ring, it will attempt to set up a new ring with itself as its only member. Such a ring is called a *self ring*. If the station trying to establish a ring finds another station willing to join the ring, the *self ring* becomes a *ring*.

The *right-to-transmit* will be passed in rotation among the ring members, with a cycle that includes every ring member. A complete transit of this ring is referred to as a *ring cycle*.

L.2.1.3. Successors and Predecessors

If station S_A is passing the right to transmit to station S_B, then station S_B shall be called the *successor* of station S_A. The *predecessor* of station S_B shall be station S_A.

In other words, the *successor* of station X is the station to which X sends the *right-to-transmit*

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

token to and the *predecessor* is the station from which X received the *right-to-transmit token*. Note that these stations are member of the same ring.

L.2.1.4.   Ring Transmit Order

The order in which the Right-To-Transmit (RTT) token is passed around in the virtual ring is called the *transmit order*, which also defines the *successor* and *predecessor* relationship among *ring members*. For example, if the RTT token is passed from station S_A to station S_B, from station S_B to station S_C, and back to station S_A, then the transmit order is $\{S_A > S_B > S_C > S_A\}$. In a ring whose transmit order equals $\{S_A > S_B > S_C > S_A\}$, station S_B is the *successor* of station S_A, station S_C is the *successor* of station S_B, and station S_A is the *successor* of station S_C.

The ring transmit order is sent to current and potential ring members as a Transmit-Order-List (TOL) which represents a directed graph of the historical transmission order. The TOL is a graph that connects a sequence of stations, which may include a station one or more times.   This graph is encoded as a list of nodes, ending with the node that transmits the TOL.

In a stable WTRP network, the TOL will represent a ring. Then *token* is passed directly between each *successor* and *predecessor* in the TOL.   In order to accommodate complex ring topologies, a node may appear more than once in the TOL. Every ring member must be included in the TOL at least once.

In a changing WTRP network, the node holding the token will choose its *successor* to optimize transmission. When this happens, the TOL will not form a ring and will represent the recent transmission history.   The TOL in this situation will generally be shorter than the previous stable ring.   Nodes will collaborate to circulate the token in order to build a new stable ring.

L.2.1.5.   Node States

The following WTRP states are defined in the protocol:

- Floating State (FLT) - the initial (starting) state where a station is not part of a ring looks for an existing ring that it can join.

- Self Ring State (SFR) - in this state a station assumes there is **no** existing ring to join, and will therefore try to setup a new ring and declare itself the ring owner. The new ring will start with this station as the only member and is therefore called a self-ring.

- Seeking State (SEK) - in this state a station considers itself to be in a self-ring and has broadcast an INVITE message as an invitation for other stations to join its ring.

- Solicit Reply State (SRP) - in this state a station tries to join another ring; it intends to respond to a received INVITE message but is waiting for a timeout to avoid congestion before sending its reply.

- Joining State (JON) - In this state a station has replied to the invitation to join (i.e., to the received INVITE message) by sending a JOIN message to the node inviting it to join the net.

- Have-Token State (HVT) - in this state a station is part of a ring. It has received a RTT (right-to-transmit) token from its predecessor and with this token, the right to transmit. From this state a

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

node may optionally transition to SLT in order to invite other nodes, after which it will return to this state.   Then it will transmit token, messages and user data and transition to MON.

•   Monitoring State (MON) - in this state a station is part of a ring. It passed the RTT token to its successor, but is unsure if the successor has received the right to transmit (i.e., the station is MONitoring the channel for an implicit acknowledgement the RTT was successfully passed);

•   Idle State (IDL) - In this state a station is part of a ring and knows it has successfully passed the RTT token its successor; it will process any D_PDU messages received from other stations;

•   Soliciting State (SLT) - in this state a station is part of a ring, currently has the right to transmit, and is inviting new stations to join the ring by broadcasting an INVITE message and listening for replies;

L.2.1.6.   Primary Timers

Most states have one or more associated timers. The primary Timers defined within the WTRP state machine are:

•   Claim Token Timer (TCLT) - Timer used in the *floating-state* (FLT). Controls the time a station waits while in the *floating state* to claim a token before transiting to another state; a station restarts its TCLT timer when it transits to the FLT state.

•   Solicit Successor Timer (TSLS) - This timer is used in the *self-ring-state* (SFR). If it times out the station shall transit to the *seeking-state* (SEK). The timer is specifed with a random timeout to reduce the probability of collisions by transmissions from stations attempting to establish different rings at the same time.

•   Solicit Reply Timer (TSRP) - Timer used in the *solicit-reply-state* (SRP). A station starts the *solicit-reply timer* when it transits to the SRP state. When it expires the station will send a JOIN message and transition to the JON state.

•   Contention Timer (TCON) - Timer used in the *joining-state* (JON). It controls the time a station waits for a response from another station following an attempt to join the network, so-named because failure to receive a response is attributed to contention with other stations attempting to join the network at the same time; a station starts its *contention timer* when it goes to JON state.

•   Idle Timer (TIDL) - Timer used in the *idle-state* (IDL). It controls the time a station waits for its *right-to-transmit* before transiting to the *floating-state* (FLT).

•   Solicit Wait Timer (TSLW) - This timer is used in the *soliciting state* (SLT) and the *seeking-state* (SEK) by stations inviting new ring members. When it expires, the inviting station will update the Transmit Order List to include all nodes that have sent JOIN messages.

•   Token Pass Timer (TPST) - Used in the *monitoring-state* (MON). It controls the time a station waits after passing an RTT (or other) token to another *station* and failing to hear an implicit acknowledgement before considering the *right-to-transmit* as lost.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.2.1.7.   *Token Contents*

The right-to-transmit-token can be considered as a simple flag that controls the right to transmit. The token is transmitted along with the identification of the "next to solicit" node. This mechanism enables nodes to take it in turns to solicit for new ring members.   This is important, as it enables growing the ring with nodes that can only communicate directly with a subset of the current ring members.

L.2.1.8.   *Promiscuous reception*

Promiscuous reception is a receive mode in which a station performs limited processing and information collection on all D_PDUs it receives, whether or not they are not addressed to it (i.e., the station takes in all traffic).

Promiscuous reception is the means by which a station discovers and confirms the existence or loss of links in the HF radio network (or that may be used to construct a network), compiles local node adjacency information, which is used to determine the TOL.   This is core to the operation of WTRP.

L.2.1.9.   *Receive Table, Connectivity Table, Adjacency Matrix and Next Hop Table*

Each node maintains a *receive table* that records information on all other ring members where transmissions have been recently received. Each ring member will transmit regularly as the token passes around the ring, so it is certain that there will be transmissions from each node in a working ring. For each node from which data has been recently received, the receive table will record that this node can be heard.   Where no data has been received from a node in the ring, the receive table will record that no data is being heard.

In addition, for each node where data can be heard, the *receive table* **shall** record a *transmit speed*. This speed reflects the maximum recommended speed for bulk throughput for data transmitted to the node.   The model and encoding is aligned to the Data Rate Selection EOW encodings specified in STANAG 5066 Annex C. This speed will be calculated from the SNR and Frame Error Rate of data received from the node.

Each node **shall** send it's *receive table* to all other nodes, by transmission around the ring whenever the receive table changes.   This means that every node will have the *receive table* for all nodes, which provides the node with information on transmission of data to the node from every other node.

The *receive table* data enables the node to build up an internal *connectivity table* for all nodes.   For each other node this table will record:
1. The connectivity status of the node. One of:
    a. No direct connectivity
    b. Bidirectional transmission
    c. Transmission only from the node to local node.
    d. Transmission only from the local node to the node.
2. Transmit speed from the node to local node
3. Transmit speed from the local node to the node

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Unidirectional transmission can be useful for non-ARQ data. For ARQ data and token transfer, it is essential that transmission is bidirectional. Because of this, primary connectivity calculations in this specification are based on bidirectional transmission, which is safe for all types of transfer. An *adjacency matrix* is built using the receive table information to show for each node pair with bidirectional connectivity the *transmission speed* in each direction.

Token transmission to nodes that are not directly connected is managed by the TOL, which provides a sequence of nodes with bidirectional data transfer between each pair.

The *adjacency matrix* and *connectivity table* can then be used to build a *next hop table* that indicates for each node in the ring one of:

1. The node can be reached directly for all transmissions; or
2. The node can be reached directly for non-ARQ transmission only; or
3. The preferred *data relay node* to be used for other communication.
4. Transmission speed.

The first two can be determined directly from the *connectivity table*, and when there is direct communication this supplies the transmission speed. For nodes that cannot be reached directly, the *adjacency matrix* can be used to determine the best *data relay node*. For each directly connected node, the adjacency matrix can be used to determine paths to nodes directly connected to these nodes. This can be done iteratively to determine the distance to each node, and a set of paths to each node.

Where several directly connected nodes can act as the *data relay node*, one of them needs to be chosen for the *next hop table*. The usual choice of the preferred *data relay node* is the one that is expected to receive the token next (which is straightforward to determine from a stable TOL), which will usually minimize transfer time. The transmission speeds of each link **may** be considered, and it may be preferable to choose a *data relay node* that leads to faster transmissions on each hop. Note that the next hop table must be recalculated whenever the TOL or *receive table* data changes.

The *adjacency matrix* is used to communicate to the DTS and CAS layers of STANAG 5066, which nodes can be accessed directly, and which nodes need to be relayed. This also implicitly identifies nodes that cannot be reached.

To send data to nodes that are not directly connected, STANAG 5066 needs to send data to the *data relay node* for that node. The *data relay node* then needs to relay the data onwards to its final destination or to the next relay hop. The protocol to do this is specified in STANAG 5066 Annex R (Routing Sublayer).

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.2.1.10. Tour

A *tour* of a graph is a sequence of nodes from the graph such that each node appears at least once and two nodes are adjacent in the sequence only if they are adjacent in the graph. An unconstrained, ordinary tour allows revisits to network nodes.

In the context of the wireless token-ring protocol, the sequence of ownership of the right-to-transmit (RTT) token, i.e., a stable Transmit-Order-List should be a *closed tour* of the network that starts and ends with nodes that are adjacent in the network, i.e., every node gets an opportunity to transmit, and passing the RTT token along the closed tour is feasible because nodes that are adjacent in the transmission sequence are adjacent in the network. A stable TOL is always a *closed tour* and may be an *unconstrained closed tour*.

L.2.1.11. Ring-Cycle Length (RCL)

The Ring-Cycle Length (RCL) is the length in hops of the tour through the WTRP network taken by the RTT token as it follows a ring transmit order list (TOL).

L.2.2. Concept of Basic Ring Operation

The basic concept of the Wireless Token Ring Protocol is to provide a mechanism that allows two or more stations operating on the same single-frequency radio channel to share it in such a way that only one station will transmit at a time. Only the station that has the *right-to-transmit* (*token-holder*) is allowed to transmit on the shared channel. The *right-to-transmit* is passed onward to all stations that joined the ring. Figure L-1 shows how the *right-to-transmit* is circulated among the *ring-members*.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4
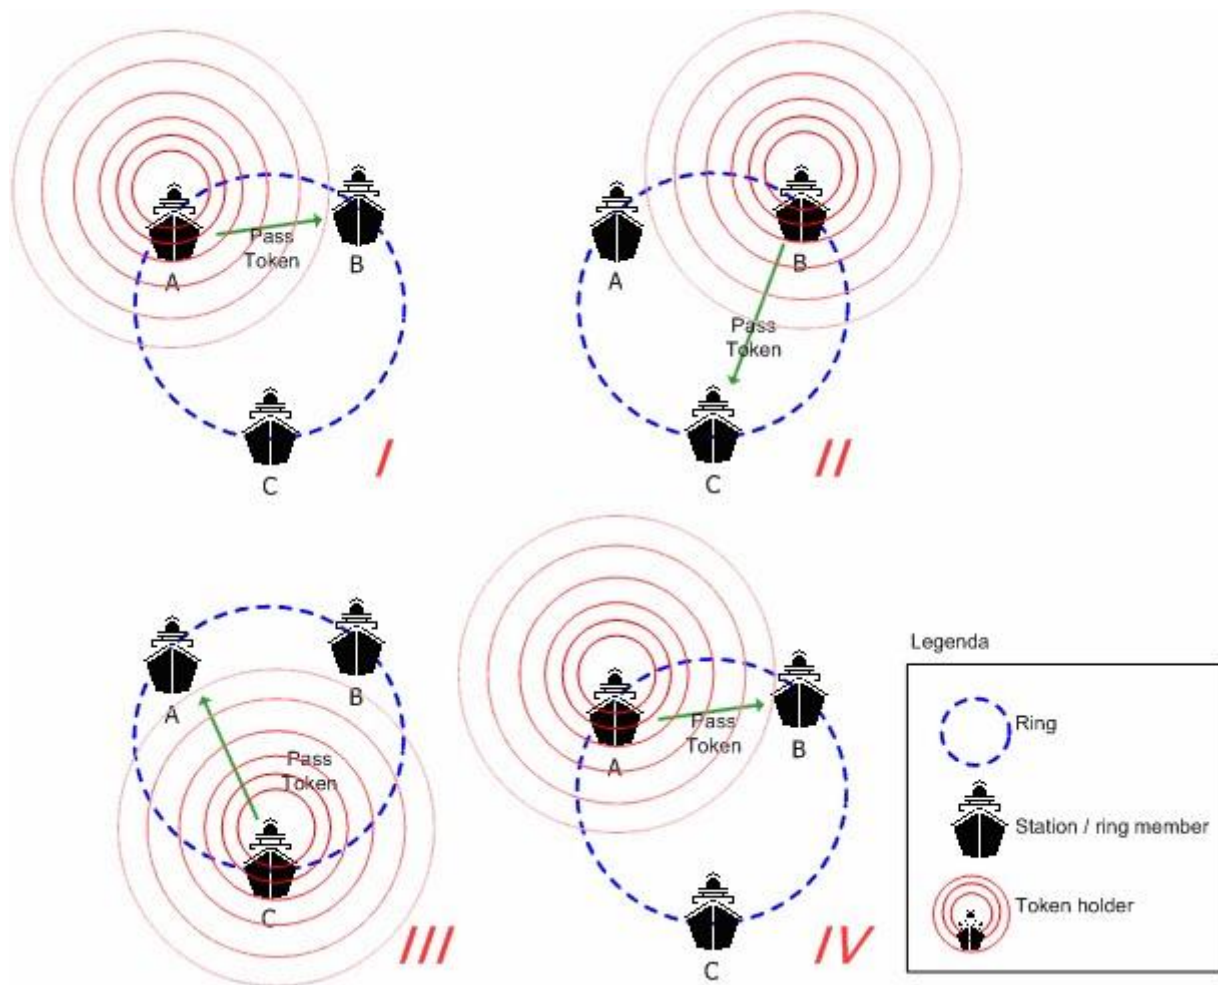
Figure L-1 - Normal Token-Ring Operation

The ring is a closed cycle of stations that transmit each in turn in a prescribed sequence, which will be adapted if new stations join the network or changes in network connectivity force adaptation. When a station receives the *right-to-transmit* from its *predecessor* it will take control of the channel. In addition to the *right-to-transmit token* the station may also receive additional WTRP messages, in particular the current *transmit-order-list (TOL)* and updated *receive table* messages from one or more nodes. These messages will not be transmitted when the information is stable. *Receive table* messages are passed around the ring to ensure that all nodes have the latest information. A node will add its own *receive table* message if this information has changed. The *right-to-transmit token* is encoded in a standard STANAG 5066 EOW message and will generally be included multiple times in a transmission to minimize risk of token loss. If the station has no data to transmit, then it **shall** pass the *right-to-transmit* and associated messages immediately.

The messages sent by a station are transferred in a set of one or more D_PDUs. The D_PDU **shall** count down the *end of transmit time* (EOT) in accordance with the requirements of Annex C. A station **shall** not exceed the maximum transmission time allowed.

Note that though the transmission sequence and ownership of the *token* in the ring is prescribed, a station with the *right-to-transmit* **may** send data to any other station in the network that it determines from the

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

connectivity table can receive data and transmit directly back acknowledgements for ARQ data, not just its successor or predecessor in the virtual ring. Any station in the virtual ring may therefore engage in multiple concurrent traffic exchanges with any other station in range. This includes the capability to support concurrent soft-link, physical-link, and ARQ connections in accordance STANAG 5066 Annex A, B, and C, with retransmission timers and other timing parameters selected to allow for the network size.

Possession of the *right-to-transmit* controls access to the radio channel; it does not prescribe the destination(s) to which traffic can be sent. The token-ring protocol controls node access to the transmission medium to preclude collisions and self-interference in the network and thereby increase the efficiency and throughput in the network. To do so, in general, there is no requirement that all nodes in the network be in communications range, only that the network have a closed tour, i.e., that a cyclic sequence of nodes exist where every node is in communications range of its predecessor and successor in the sequence. To support WTRP networks where direct communication is not possible between all nodes, the WTRP layer communicates connectivity to the higher layers of STANAG 5066.   The higher layers will only send data to nodes that can be reached directly.   Relay to other nodes is provided by STANAG 5066 Annex R (Routing Sublayer).

Each station will retransmit the *token,* passing the *right-to-transmit* to its *successor* until this has been acknowledged by its *successor.* This acknowledgement is implicit, by observing when the successor transmits data.

L.2.3.　　　　　　Ring formation

Before operating as a ring, a ring must be formed by stations on the same radio channel. The ring is a self-organizing and self-repairing mechanism, subject to a number of requirements:

- As soon as a station starts to operate it will listen to the radio channel in a floating state, unaffiliated with any ring.

- A station may transmit data (i.e., D_PDU other than ring-management messages) only if it is part of a ring that is not a self ring.

- A station shall first listen for an existing ring on the radio-channel; if it hears a ring it will try to join that ring, otherwise it will form a new ring (a self ring).

- If a station receives the right-to-transmit from another node, it has become a part of the ring.

- After initial receipt of the token, the node will pass the token back to the node that gave it the token.

- A station dropped from a ring must re-join the ring to become again part of it and obtain transmission rights.

There are two possible ways to form a ring:

1. Join an existing one.

2. Start a new ring yourself.

The process of Ring Formation is illustrated below in Figure L-2 - Ring Formation.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Listening: Trying to find an existing ring.	Start "self-ring" waiting for a random timeout to broadcast message.INVITE to join the ring	TSLS-timeout: Broadcast message.INVITE	Listen for stations which will reply to join the ring with a message.JOIN	token received	Waiting for SEK state to time out (TSLW)	Ring setup complete: start passing token

Set time-out for random timeslot (TSRP) to reduce collisions

TSRP timed out. replying to SLS with message.JOIN.

Figure L-2 - Ring Formation

Figure L-2 shows a time-sequence evolution of two stations, station A and station B, as they form a two-node ring.

In situation *I* station A is listening for an existing ring. After its TCLT timer times out, station A assumes there is no existing ring (otherwise it would have heard it) and it transits from a FLT state to the SFR state, forming a *self ring* as illustrated in situation *II*.

The transition moment from the SFR state in situation *II* to the SEK state in situation *III* is controlled by a random timeout that avoids collision between two stations in the same state. The value of the timeout is calculated based on picking an asynchronous timeslot (See TSLS timer details in section L.5).

In the SEK state of situation *III* a station invites new members by sending an INVITE message. Each station that is seeking to join will reply with a JOIN (join ring) message. There could be more than one station waiting to reply to an INVITE message. Therefore, to reduce the probability of collision between replies from multiple joining nodes, the joining node's reply shall be transmitted in a randomly chosen time slot among a set of slots available for replies. The station that sent the INVITE message will add all of the nodes requesting to join to the Ring by passing the token to each of them in turn, noting that on initial token receipt a node **shall** pass the token back to the node that gave it the token. All the new nodes will now be included in the TOL, which records token transmission history, and thus will have become members of the ring.

Once the ring is created, each node periodically creates an opportunity to invite new stations to join the ring by broadcasting an INVITE message and waiting for JOIN message replies from nodes wishing to join the ring.

L.2.4. Ring Entry

The invitation to join shall be made periodically by a station receiving the right-to-transmit, with invitation frequency about once every SLS_INTERVAL seconds. The opportunity to invite new members

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

needs to be shared between ring members, as new nodes may only be able to connect to one existing member. The *token* identifies the next node to invite new members. Once this node has issued an invite, it will pass the invite to another member of the ring, by setting the value in the token. This is done using a mechanism which gives the invite opportunity equally to all ring members. A configurable maximum ring size **may** be specified, which can be used to stop rings growing too large. This maximum ring size **may** also be used when the number of nodes is known, to prevent taking time to invite new members when there are no possible new members. Invitations can also be configured so that they are less frequent (lower overhead) for stable rings.
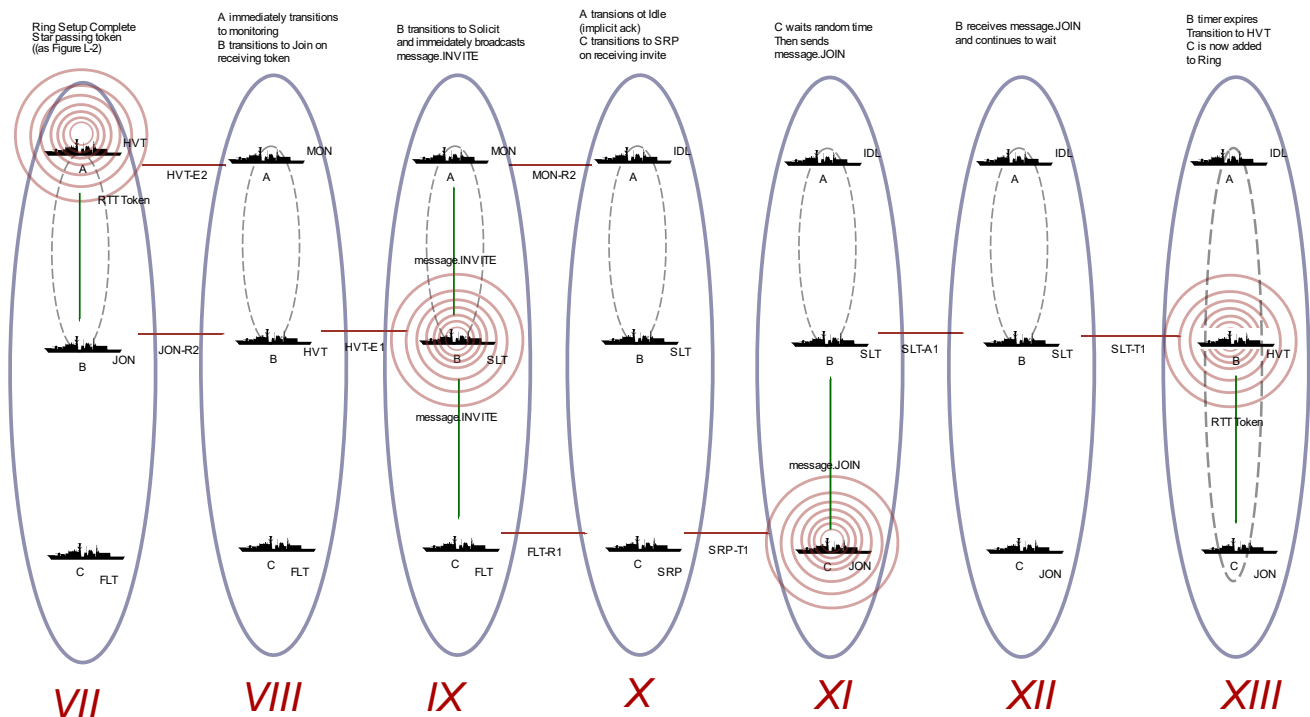
Figure L-3 - Ring Entry

Figure L-3 illustrates adding a station to an existing ring and continues the example of Figure L-2, showing addition of station C to the ring.

In VII, station B has received the token as in HVT state. Station B determines that it needs to issue an invite, so immediately transitions to SLT (IX) and immediately broadcasts a message.INVITE. Station A receives the message.INVITE, which it uses as implicit acknowledgement of token transfer, so it transitions from MON to IDL (X). Node C in FLT plans to accept the invitation and so transitions to SRP (X).

Station C waits for a random time before responding with message.JOIN (XI). Station B receives the message.JOIN (XII) and continues to wait for potential other message.JOIN from other stations. After a timer, station B transitions to HVT and station C is added to the ring. Station B sends message.JOIN to station C (VIII).

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.2.5. Connectivity Update

Each node monitors the connectivity from all other ring members. It **may** also monitor connectivity from other nodes, which is recommended as it will provide useful information in the event of a monitored node joining the ring. This information from other ring members is recorded in the *receive table*, which notes nodes that are heard and an associated recommended maximum transmit speed, determined from SNR, FER and other information.

Nodes will be removed from the table on either reaching a configurable maximum age or a configurable number of ring cycles with no data heard.

L.2.6. Communicating Connectivity

The calculated receive table is shared with all other nodes. This is done by sending an TABLE message that contains the receive table. TABLE messages are sent along with the token and will progress around the ring. A node sending an TABLE message will ensure that it arrives with the token after a complete ring cycle (noting that a node may get the token multiple times in a ring cycle). When this has happened, the node will know that all other nodes in the ring have the current receive table and so there is no need to transmit the receive table again.

Whenever the receive table changes, it is sent around the ring to share this information. In this way, all nodes have full connectivity information.

L.2.7. Stable Rings

The TOL provides immediate history of the token transfer. When a TOL is received, the first node in the TOL will always be the predecessor of the node receiving the TOL. The nodes preceding that represent token transfer history.

When the ring is stable, the TOL will represent a closed tour, and the last node in a received TOL will always be the local node. The TOL is only transferred if it has changed since the last transfer between a predecessor/successor pair. Both nodes know the previously transferred TOL, so it is only transferred on change.

When received TOL is transmitted to the successor in a stable ring, the local node is removed from the back of the TOL and added to the front of the TOL. The successor will now be the last node in the TOL.

L.2.8. Changing Ring Connectivity

There are four scenarios where the TOL will change.

- Scenario 1 (Joining Scenario): A node joins the network and will be added to the TOL. Every node member of the the ring shall be in the TOL;

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

- Scenario 2 (Leaving Scenario): A node chooses to drop out of the ring, either due to operator request or due to detecting another ring. In the latter case, the whole ring is closed, with the expectation that the nodes will join the new ring. A node leaving the ring is removed from the TOL before the TOL is transmitted.

- Scenario 3 (Token Transfer Fail Scenario): Changes in the network may lead to a situation where token transfer in a stable ring to the current successor node fails. This node will then choose a new successor to transfer the token to. This **may** lead to nodes being dropped from the ring.

- Scenario 4 (Ring-Optimization Scenario): It is possible that in a stable ring, that a node can determine a better order, with a different successor to the one in the current stable ring. To move to this new order, a node will transfer the token to this new successor.

These four scenarios are considered in more detail below.

After these changes, the immediate TOL will represent transfer history, which may well not be a ring. Nodes will work co-operatively to move from this state towards a new stable ring. In this situation, the TOL performs two key functions:

1. It communicates the current list of ring members.
2. It shows the node which was visited longest ago. In choosing rings, it will be desirable to visit this node as soon as possible.

There is no benefit for the TOL being longer than necessary to achieve these functions. Therefore, a transmitted TOL will be truncated so that it ends in the last unique node. Where it is determined that a node at the back of the TOL cannot be reached, it is removed from the TOL.

L.2.8.1.   *Joining Scenario*

A new node joining the network results in a new network topology.

The inviting node broadcasts an INVITE message, to which one or more joining nodes respond with a JOIN message. The JOIN message contains the receive table for the joining node, which will help the inviting node select transmit speed to that node. The inviting node will pass the token in turn to each node that is being added to the ring. When a joining node received the token for the first time, it **shall** pass the token straight back to the inviting node.

It can be seen that that this process will add each new node into the TOL (transfer history). Ring optimization will often follow joining.

L.2.8.2.   *Leaving Scenario*

A node choosing to leave a ring **shall** do so by waiting until its transmit opportunity and then forwarding the right-to-transmit to its successor, including monitoring for correct handover. The node **shall** remove itself from the TOL so that it is not included in the transmit history.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.2.8.3.   *Token Transfer Fail*

Instances will occur with topology change where a node holding the RTT token in a stable ring fails to transfer the token to the designated successor.   When this happens, the node will determine an alternate successor and transmit the token to this successor.   This will break the stable ring, and the next node will follow the ring optimization process.

L.2.8.4.   *Ring Optimization*

When the TOL forms a ring, the node in the ring may determine that there is a different ring order that will give better performance **and** that this ring needs a different choice of successor to the current one (the last node in the TOL).   In this situation, the node will transfer the token to the successor determined from the new ring.   The transmitted TOL will be truncated (if necessary) so that it ends with the last unique node in the ring. A node **shall** always put itself at the front on the transmitted TOL.

L.2.8.5.   *Moving Towards a Stable Ring*

After all of the above scenarios the transmitted TOL will generally not form a ring. A receiving node can detect a TOL which is not a ring, because the local node will always be the last entry of a received TOL.   In this scenario, the node **shall** determine an optimal ring and transmit the token to the successor determined by this ring.

Nodes will work co-operatively to build a stable ring, noting that different implementations **may** determine different rings as optimal.   Two strategies are mandated in order to prevent interactions between such strategies preventing a stable ring from forming:

1. Give priority to reaching the last node in the TOL.   This will help ensure that nodes do not get left out.
2. Avoid changing ring topology for marginal performance gain.

L.2.9.   **Service Provision to Higher STANAG 5066 Layers**

STANAG 5066 Annex J defines the MAC model which WTRP provides.   There is a clean layer model, with M_ service elements communicating between the layers.   The details of these service elements are not specified, but the model and associated protocols are clear.   WTRP provides some additional services to the generic MAC functions defined in Annex J.

L.2.9.1.   *Data Rate, Interleaver and Transmission Length*

STANAG 5066 DTS layer will control selection of transmission speed, interleaver and transmission length. The DTS layer will also determine which D_PDUs to duplicate, such as sending ACKs multiple times to improve reliability.

WTRP introduces new D_PDUs with rules for transmission and duplication.   Selection of data rate, interleaver and transmission length must be cognizant of this information.   WTRP also replaces the Annex C

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Data Rate Selection mechanism with a simpler service that provides recommended maximum transmission speed for each directly connected node, using the *connectivity table* defined in L.2.1.9.

L.2.9.2. Node Availability and Routing

WTRP determines which nodes in the ring can accessed directly and for those nodes which cannot be accessed directly it specified a preferred relay node. This information is generated in the *next hop table*, as specified in L.2.1.9. This information is passed upwards as an Annex J M_ service primitive. This information can be used:

1. By the CAS as specified in STANAG 5066 Annex B to immediately reject messages to nodes that are not in the ring; and
2. By the Routing Sublayer as specified in STANAG 5066 Annex R to relay data to nodes that are not directly connected.

L.2.9.3. Broadcast Retransmission

When a broadcast/multicast PDU is received, information is provided to the higher layers. It is important that broadcast and multicast work correctly for nodes that are not directly connected. This needs to be done in a way that prevents duplicate onward broadcast. Consider four nodes connected in a long trapezoid shape, where the two end nodes cannot communicate directly. If one of the end nodes sends a broadcast message, then one (but not both) of the middle nodes that can communicate with both nodes needs to broadcast the message, so that all nodes receive it.

This is achieved by providing a *Relay Responsible List* to the higher layer for any broadcast or multicast PDU that is received. This tells the higher layers which nodes it needs to re-broadcast to. Handling this is described in STANAG 5066 Annex R (Routing Sublayer).

To do this calculation, the node needs to handle connectivity data as if it is the node which sent the broadcast. The set of nodes which will have received the message directly can be determined. The additional set of nodes that need to receive the message can be determined from this list and from the membership of the broadcast list or multicast list (noting that the node knows the identity of all nodes). For each of these nodes, the preferred relay node (from the perspective of the sending node) can be determined. The *relay responsible list* is the list of nodes for which the preferred relay node is the local node.

L.3. WTRP TOKEN AND MESSSAGE SPECIFICATION

L.3.1. Node ID

Nodes are identified by a single byte ID (range 0-255). This limits the number of nodes in a ring to 256, which is expected to be higher than the practical limit. As nodes need to be extensively referenced in the protocol, a compact representation is important. Externally, each node is identified by a variable length STANAG 5066 address (typically 3.5 bytes).

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

The Node ID will usually be the same as the last byte of the node's STANAG 5066 address, unless there is a conflict in the ring.

Each Node will maintain a map between Node ID and STANAG 5066 address. The mapping is derived from the set of receive tables; when distributed the receive table includes this mapping. A node must have a complete set of receive tables, with a receive table for all nodes in the ring. The WTRP protocol enables "missing" receive tables to be requested.

This mapping is assumed in a number of the procedures below, so that nodes can be identified either by Node ID or STANAG 5066 address.

L.3.2. Token Protocol

The token is transmitted using a standard EOW message. This enables the token to be transmitted along with standard data. The EOW must be sent in a message with destination address of the token recipient. The source address will always be that of the token sender.

The single byte value of the EOW is set to the node ID of the node that currently has responsibility to invite new nodes. This ensures robust transfer of this information along with the token.

Two types of EOW are used.

1. Token Only. This uses EOW value 13. This message transfers token only. The TOL is unchanged since last token transfer between this pair of nodes, and so does not need to be transferred. There may be messages containing new receive tables. Loss of receive table is non-critical, as receive table updates will be repeated if they are not received back at the node originating the update.

2. Token and TOL. This uses EOW value 14. When this EOW is used the TOL has changed since last token transfer between this pair of nodes and the TOL **shall** be transferred. In the event that the TOL is not received, the receiving node **shall** request retransmission of the TOL using a RETRANS message. The receiving node **may** also request retransmission of TABLE messages, which may be necessary to map from unique IDs in the TOL to STANAG 5066 addresses.

L.3.3. WTRP Messages

WTRP uses a set of messages to communicate information. These use EXTENSION D_PDUs as specified in STANAG 5066 Annex C. The allowed messages and defined extension D_PDU numbers are specified in Table L-1:

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Message	Description	Extension Number
TOL	Transmit Order List	1
TOL+	Transmit Order List (Extended Form)	2
TABLE	Receive Table	3
TABLE+	Receive Table (Extended Form)	4
INVITE	Invite other nodes to join ring	5
JOIN	Request to join ring	6
RETRANS	Retransmission Request	7
CLOSE	Close the current ring	8

Table L-1 – WTRP Messages

The syntax of these messages is specified in the following sections, and procedures to use them are specified in Section L.4.

A number of the messages will often be short enough to fit within a STANAG 5066 header, but will not always be. Because of this, two forms of the TOL and TABLE message are defined.   Use of different extension number allows the format to be determined from the extension number in the D_PDU header.

L.3.3.1.   TOL Message

	MSB							LSB
	7	6	5	4	3	2	1	0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)				Size of Header (h)			
3+m	Source and Destination Address							
4+m	MSB Extended D_PDU Type = 1							LSB
5+m	TOL (max 24 bytes)							
h+m-2								
CRC	CRC on Header							
CRC								

Figure L-4– TOL Message

L-17

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

	MSB							LSB
	7	6	5	4	3	2	1	0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m	MSB Extended D_PDU Type = 2							LSB
5+m	MSB Length of TOL (n)							LSB
CRC	CRC on Header							
CRC								
h+m+1	TOL							
h+m+n								
h+m+n+1	MSB							
	CRC							
h+m+n+4								LSB

Figure L-5– TOL Message (Extended Form)

The TOL Message communicates a TOL.   Each byte in the TOL represents a node identified by its node ID. The first byte is the oldest node in the TOL and the last byte is the most recent node (which will always be the node transmitting the TOL).

The TOL Message encodings have the standard S55066-EP10 fields shown without colour. The extended D_PDU type is set to 1 for TOL and 2 for TOL (Extended Form).   In the extended form, the Length of TOL header field specifies the length of the TOL.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.3.3.2 TABLE Message

	MSB							LSB
	7	6	5	4	3	2	1	0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m	MSB Extended D_PDU Type = 3							LSB
5+m	MSB Sending Node Address							
8+m								LSB
9+m	MSB Sending Node ID							LSB
10+m	MSB Version							LSB
12+m	Table (max 18 bytes)							
H+m-2								
CRC	CRC on Header							
CRC								

Figure L-6– TABLE Message

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

	MSB 7	6	5	4	3	2	1	LSB 0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)				Size of Header (h)			
3+m	Source and Destination Address							
4+m	MSB Extended D_PDU Type = 4							LSB
5+m	MSB			Sending Node Address				
8+m								LSB
9+m	MSB			Sending Node ID				LSB
10+m	MSB			Version				LSB
11+m	MSB			Number of Table Entries (n)				LSB
CRC CRC	CRC on Header							
h+m+1	Table							
h+m+n*2+1	MSB							
				CRC				
h+m+n*2+4								LSB

Figure L-7– TABLE Message (Extended Form)

The TABLE Message communicates a receive table.   The Table is a sequence of two byte entries in arbitrary order.   Tables of nine entries or less **shall** use the TABLE Message shown in Figure L-6. Longer tables **shall** use the TABLE Message (Extended Form) shown in Figure L-7.   The fields are the same, except that the Extended Form has a Number of Table Entries header, which enables the length of the table, and thus the D_PDU length to be determined.

The Sending Node Address field is the STANAG 5066 address of the node that generated the receive table. The Sending Node field is the node ID of the node that generated the receive table.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

The table has a version number starting at 0, which is incremented for each change, resetting to zero after 255. This is to ensure that when multiple versions of a receive table are circulating, that a node will ignore a version prior to the one it holds.

Each table entry comprises two bytes. The first byte is the node ID of the node that data can be received from. This **must** be a node in the TOL. The second byte is the recommended transmit speed. This is encoded following the encoding specified in Section C.6.5.1 of STANAG 5066 Annex C (Transmission Speed and Interleaver Encoding).

L.3.3.3 INVITE Message

	MSB							LSB
	7	6	5	4	3	2	1	0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)				Size of Header (h)			
3+m	Source and Destination Address							
4+m	MSB Extended D_PDU Type = 5							LSB
CRC	CRC on Header							
CRC								

Figure L-8 – INVITE Message

The INVITE message is used by a node with the token to invite other nodes to join. It is sent to the broadcast address, so that any listening node will receive it.

L.3.3.4 JOIN Message

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

	MSB							LSB
	7	6	5	4	3	2	1	0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m	MSB Extended D_PDU Type = 6							LSB
5+m	Table (2 bytes)							
6+m								
CRC	CRC on Header							
CRC								
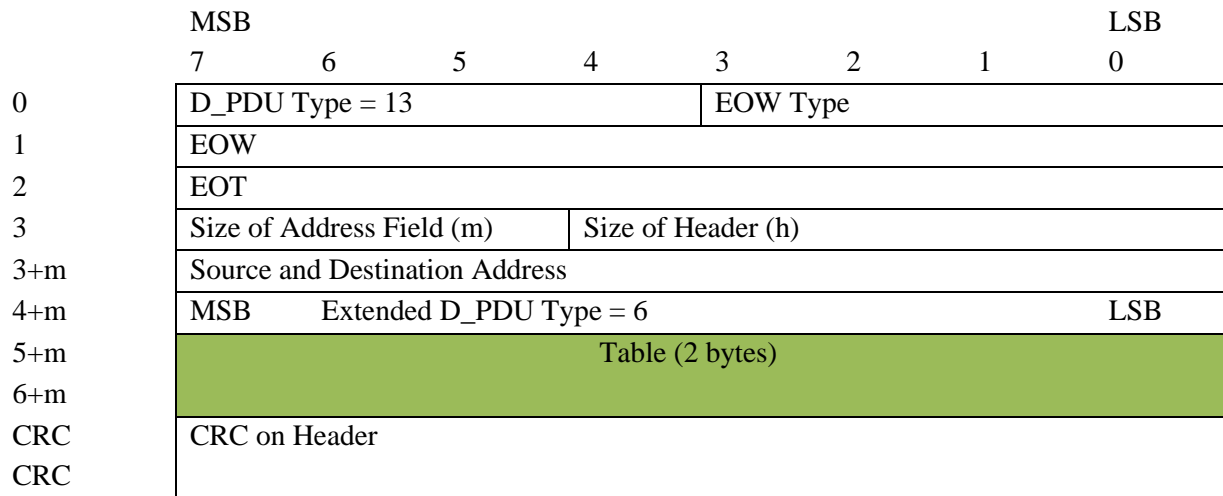
Figure L-9– JOIN Message

The JOIN message is used by a node responding to an INVITE message, indicating that it wishes to join the ring.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.3.3.5 RETRANS Message

	MSB							LSB
	7	6	5	4	3	2	1	0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m	MSB Extended D_PDU Type = 7							LSB
5+m	Missing Receive Table List (0-24)							
CRC	CRC on Header							
CRC								
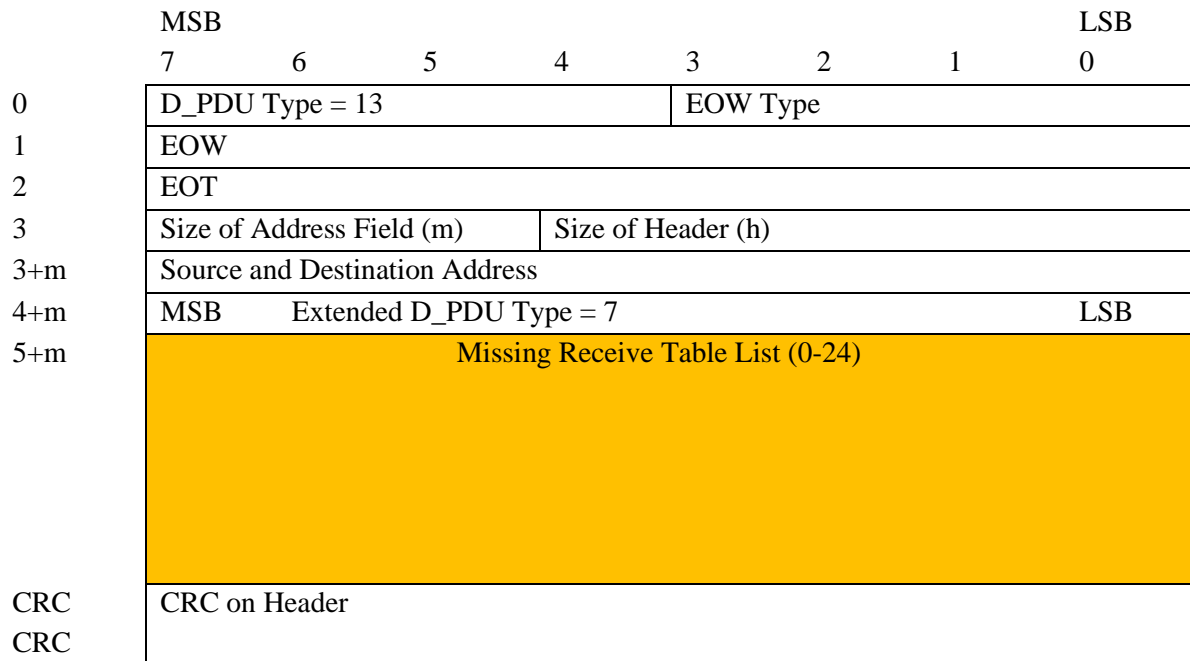
Figure L-10 – RETRANS Message

The RETRANS Message is used when a node receives a token but is missing critical information. There are two scenarios where this can happen:

1. Where the Token is received with an EOW that indicates that a TOL was sent and no TOL message is received. In this case, the message is sent without a Missing Receive Table List. This indicates a request to resend the TOL.

2. Where a TOL is received, but the receiver does not have a receive table entry from one or more of the nodes in the TOL. Each byte of the Missing Receive Table List is the Node ID of a required receive table.   The RETRANS message can encode up to 24 nodes.   If more receive tables are needed, multiple RETRANS messages are used.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.3.3.6 CLOSE Message

	MSB							LSB
	7	6	5	4	3	2	1	0
0	D_PDU Type = 13				EOW Type			
1	EOW							
2	EOT							
3	Size of Address Field (m)			Size of Header (h)				
3+m	Source and Destination Address							
4+m	MSB Extended D_PDU Type = 8							LSB
CRC	CRC on Header							
CRC								

Figure L-11– CLOSE Message

The CLOSE Message is used to signal to close a ring.   This will be used when a node has detected another ring, which it intends to join after leaving the ring, with the expectation that nodes in the ring being closed will subsequently join the new ring.

L.4.   PROCEDURE OF OPERATION

WTRP operation is defined in terms of the state machine specified here. Specifications and definitions are provided for states and state transitions.   This makes use of a number of data structures and procedures defined in this section.

L.4.1.   Node Data Structures

A node shall maintain the data structure listed Table L-6.   These data structures are all prefixed with the following prefixes to facilitate clarity when they are referenced:

- "Node.":   Generic information associated with the node.
- "TOL.":   Information associated with the TOL.
- "LocalTable": Information associated with the local receive table.
- "PeerTables.":   Information associated with receive tables from other nodes. .
- "Invite.": Information associated with control of invitations.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Table L-2 – Node Data Structures

Register / Flag	Type	Initial Value	Description
Node.address	STANAG 5066 Address	n/a	The station's address. The address is expected to be configured.
Node.ID	Integer	Last byte of Node.address	Node ID of the station. May be re-assigned when the node joins a ring.
Node.cycleCount	Integer	0	The number of times that a complete ring cycle has been completed. Note that a cycle is only considered complete when the TOL has not changed for the cycle. This is to help ensure information is propagated to all nodes. Initalized to zero.
Node.Predecessor	STANAG 5066 Address	n/a	Set to sender of Token, when token is received.\n\nThis is needed when a node appears multiple times in the TOL, to identify the current TOL position.
Node.Sucessor	STANAG 5066 Address	*n/a*	Set to the Token receiver, when token is transmitted\n\nThis is needed for monitoring after the token has been passed.
Node.TokenTransferred	Boolean	false	Acknowledgement is implicit, by monitoring traffic. This variable is set if traffic arrives that indicates token transfer
Node.CloseRing	Boolean	false	Set if node detects another ring. This will lead to the ring being closed down.
Node.OperatorDrop	Boolean	false	This variable is set by an operator. If set, it will cause the node to be dropped from the ring.
Node.txTokenCounter	Integer	0	Used to record repeats of token transmission.
Node.NextHopTable	Next Hop Table	empty	This is maintained so that higher layers of STANAG 5066 can select best speed for sending to a given node and can send indirectly to nodes that cannot be reached.
Node.NewSuccessor	Boolean	False	Use in MON state to control switching successor

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Node.NewNodes	List of Nodes	Empty	Used to record JOIN requests, to build list of nodes to be added to ring
TOL.Received	Array of Received TOLs, indexed by Node ID	empty	Records received TOLs, so that TOL can be determined with a Token Only EOW is received
TOL.Transmitted	Array of Transmitted TOLs, indexed by Node ID	empty	Records transmitted TOLs, so that if TOL is the same as previous transmission that a Token Only EOW can be used
TOL.Current	TOL	Empty	Set to received TOL and updated to transmitted TOL after transmission is acknowledged
TOL.CurrentTransmitted	TOL	empty	The current TOL that has been transmitted to successor. It will be used to update TOL.Transmitted and TOL.Current when the token transfer is acknowledged.
LocalTable.Table	Receive Table	Empty	The value of the local receive table, holding information on receive quality from peer nodes
LocalTable.TableChanged	Boolean	false	True if receive table has been changed since circulation of the receive table started.
LocalTable.Circulated	Boolean	false	A table is being circulated
LocalTable.CiruclatedCount	Integer	0	The ring cycle count when the receive table was circulated. This is used to ensure circulation to all nodes
LocalTable.UpdateTimes	Array of Times, indexed by node ID	empty	For each node in the TOL, the time that the entry in LocalTable.Table was last updated.
PeerTables.Tables	Receive Table List	empty	The active list of receive tables from peers. This list is updated whenever a more recent receive table is received from any node.
PeerTables.ReceivedDirect	Received Table List	empty	List of receive tables from previous node. Need to relay on, only values received directly from predecessor. This ensures that if originator gets back the table it sent, that all nodes have a copy of the table
PeerTables.MissingList	Set of Node IDs	Empty	This is set when a TOL is analysed and there is not a receive table for one or more ring

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

			members. This gives a list of nodes for which to request receive table retransmission.
Invite.LastTime	Time	Current time	When the last invite was sent by current node
Invite.LastCycleCount	Integer	0	Ring Cycle Count when invitation last sent.
Invite.Number	Integer	0	Number of invitations issued since joining ring
Invite.Next	Node ID	0	The node that should issue invitations next, and will be sent with each token

L.4.2. Message and Token Notation

The notation specified in this section is used to refer to tokens and messages and their components. Tokens have the following references, which are used to describe both reading and setting the token.

Table L-3 – Token Notation

Notation	Type	Description
token.sender	STANAG 5066 Address	Node sending the token from D_PDU source address
token.receiver	STANAG 5066 Address	Node to which the token is being sent taken from D_PDU destination address
token.newTOL	Boolean	Set to true if a TOL message is being sent in the same transmission as the token. Retransmission or the TOL **must** be requested if it is not received. This value is set from the EOW type.
token.inviterID	NodeID	The node ID of the node which is responsible for sending the next invite and for updating this field to set the subsequent inviter.

Messages are reference by notation of the form message.<message name>, for example message.TOL. The following references to components of messages are used.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Table L-4 – Message Notation

Notation	Type	Description
message.<message name>.sender	STANAG 5066 Address	Node sending the message from D_PDU source address
message.<message name>.receiver	STANAG 5066 Address	Node to which the token is being sent taken from D_PDU destination address
message.TOL.TOL	Transmit Order List	TOL in a TOL message. When set, the RCL (Ring Cycle Length) element is set from the TOL length
message.TABLE.Table	Receive Table	Includes sender node ID and length
message.INVITE.TOL	Transmit Order List	As for TOL
message.JOIN.Table	Receive Table	As for TABLE
message.RETRANS.TOL	Boolean	Request to retransmit TOL
message.RETRANS.MAP	Boolean	Request to retransmit MAP

L.4.3. Node Procedures

The following functions are used in the state machine to specify node behaviour. Note that while this functional notation specified behaviour, this specification imposes no requirement on an implementation to use these functions.

InitializeNode()

    Return Type: None

    Description: Initializes node.   The following values are set:

- TOL is initialized to a single entry for the local node
- Other values are set to the default values as specified in Table L-6

isMember(<node address>)
    Return type: Boolean

    Description: Returns true if the given *node address* is a ring member, otherwise false.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Inspect TOL.Current to determine if the specified node is present.

RingMemberCount()

Return type: Integer

Description: Returns the number of unique nodes in the TOL.Current.

RCL ()

Return type: Integer

Description: Returns the number of nodes in the TOL.Current (Ring Cycle Length).

BreakLink(<node address>)
Return type: None
Description: When a node cannot be reached, this is used to change network topology. This is used to recalculate the TOL to determine a new successor.

This is achieved by finding the entry for <node address> in PeerTables.Tables and removing the entry for the local node in the identified receive table.   This changes the connectivity record, to indicate that the node in question cannot be reached from the local node.   The TOL can then be optimized to identify a node that may be reached.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

RetransmitMessages(<message.RETRANS>)

      Return Type: None

      Description: Used when message.RETRANS is received

If Missing Receive Table List in message.RETRANS.TOL is empty, retransmit the TOL in message.TOL.

If Missing Receive Table List in message.RETRANS.TOL is not empty, retransmit each of the requested receive tables in a message.TABLE message for each receive table.

The token **must** also be transmitted following L.4.8.   Other WTRP messages **may** be retransmitted following L.4.8. which is recommended. It is recommended to transmit at conservative speed and to send the requests messages multiple times in the transmission.   User data **must not** be sent.

L.4.4. Processing Inbound Transmissions

This section, in conjunction with the following sections (L.4.5 – L.4.7) describes the process for handling inbound transmissions. This functionality is driven from the state machine by a single call of Receive(). This will listen for a call and continue processing until a full transmission has been received.   At the end of transmission it will return EOT and optionally one of the following as events to the state machine:

1. EOT Event.   To indicate end of transmission This is always returned, and the following associated Booleans are set.
    a. D_PDUs received.   Set to true if any valid D_PDUs received.
    b. Joins received. Set to true if Message.JOIN directed to another node is received
2. One of the following may also be returned with the EOT event:
    a. Token. If a token directed to the local node has been received and none of the following messages.
    b. Message.INVITE. If this has been received, it will be the only message and no Token will be received.
    c. Message.JOIN.   If this has been received, it will be the only message and no Token will be received.
    d. Messsage.RETRANS with Token. If this message been received, Token will always be received.

Processing of Message.TABLE and Message.TOL received is handled by this procedure and transparent to the state machine.

When a transmission is received it is fully processed until the EOT, prior to any actions being taken in the state machine. D_PDUs other than WTRP messages are processed following the rules of STANAG 5066.   If no other event is returned, EOT is returned at end of transmission.

WTRP Message.INVITE is always sent to the broadcast address.   A transmission with a message.INVITE may contain multiple copies of this message.   Message.INVITE is returned

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

WRTP Message.RETRANS is sent to a single address and will not have any other messages, but may have duplicates and **shall** include the token.   If this is received, Message.RETRANS with Token is returned.

WRTP Message.JOIN is sent to a single address.   A single transmission may contain multiple copies of this message.

If Message.JOIN.receiver is not the local node, it reflects a node that is responding to an invitation from a different node.   This node may or may not be added to the ring.   This message is ignored and the procedure finishes.

If Message.JOIN.receiver is the local node, this is the response to an invitation.   This is processed by this procedure, by adding Message.JOIN.sender to the Node.NewNodes list.   The procedure then finishes.

The sender of the transmission can be identified from any D_PDU in the transmission.   Use of isRingMember(sender) determines if the sender is in the current ring.   If the sender is not a ring member, the procedure of L.4.10 is followed.

Once a transmission has been processed, it will be possible to determine either the target destination node for the WTRP information, or that the message is a broadcast message.INVITE.

If the target destination node is the local node, the Section L.4.6 is followed. If the target destination is another node then Section L.4.5 is followed.

L.4.5.   Handing Transmissions Directed to Other Nodes

Transmissions directed to other nodes are not handled by the state machine. However, WTRP must listen for these transmissions (promiscuous mode).   Information in these transmissions is used to update local information.

Message.JOIN, Message.INVITE and Message.RETRANS are special transmissions with no user data. Handling these is covered in L.4.5.

A message containing a token directed to another node indicates explicitly that the transmission is directed at another node.   WTRP messages will always contain a token, so it will be always be possible to   determine where a transmission is directed when it contains WTRP messages.

Arriving tokens may be noted.   Token.sender and token.receiver may be useful to provide as operator information to monitor progress of the token around the ring.

If the transmission sender of any D_PDU is Node.Successor, set Node.TokenTransferred to true.

If an arriving D_PDU contains the WTRP token and   isMember(token.Sender) is true, set Node.TokenTransferred to true.   This setting is used to change out of MON state.

Received Message.TABLE messages are used to update local status. If the message.TABLE sender node ID is the local node, ignore. If the receive table is from the current ring, update the table in PeerTables.Tables with message.TABLE if the version is more recent.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Message.TABLE has version numbers encoded as a single byte.   To compare version numbers of current and new,   Mod(new – current, 256) < 127 will be true if new is more recent than current.

L.4.6. Procedure for Receiving the Token

This section describes how to handle a message with a token that is directed to the local node and therefore needs processing by the local node.

Invite.Next is set to token.inviterID.

If token.newTOL is false, the token from last time is used.   Set TOL.Current to TOL.Received[token.sender].

If token.newTOL is true and there is a Message.TOL in the transmission, send TOL.Current to Message.TOL and sent TOL.Receiveed[token.sender] to Message.TOL.

If token.newTOL is true and there is no Message.TOL in the transmission, set TOL.Missing to true.   TOL retransmission will then be requested using message.RETRANS and no further processing is done.   Other WTRP messages received are retained for processing once the TOL is received.

The next step is to update PeerTables, and in the event that any entries are missing, to request them.

PeerTables.ReceivedDirect is cleared, unless a retransmission of receive tables has been requested.

Message.TABLE messages for nodes other than the local node are considered. If the Message.TABLE node ID does not identify a ring member, the Message.TABLE ignored.   For receive tables from ring members both of the following actions are taken for each Message.TABLE:

1. Compare message.TABLE with the value for the node in PeerTables.Tables. If the version in message.TABLE is more recent, update PeerTables.Tables with this receive table.
2. Add the most recent version of this receive table to PeerTables.ReceivedDirect.   This enables passing on the list of messages received, so that updated receive tables circulate back to the node that generated the update.

For each node in the TOL, check that there is a receive table entry in PeerTables.Tables.   If there are any missing, set PeerTables.MissingList to the Node IDs of the missing entries.   This will lead to requesting the missing receive tables using Message.RETRANS.   No further processing is done.   Retain any Message.TABLE for the current node for processing once the missing receive tables are received.

If message.CLOSE is received, set Node.CloseRing to true.

For each entry in LocalTable.Table look at the associated update timestamp in Local.TableUpdateTimes.   If it is older than RECEIVE_TABLE_EXPIRY_AGE, remove the entry from LocalTable.Table and set LocalTable.TableChanged to true.

If LocalTable.TableChanged is true, clear LocalTable.Ciruclated.   As a new receive table will be sent, there is no need to check if circulation of the previous one completed.

Message.TABLESs must be circulated around the ring, and two checks need to be made to ensure this has happened.   If the message has correctly circulated around the ring, then mark so that the check is turned off. If the message has not correctly circulated around the ring, force it to be circulated again. The checks are:

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

1. If LocalTable.Circulated is true and no message.TABLE received for the local node, set LocalTable.TableChanged to true to force repeat circulation.
2. If LocalTable.Circulated is true and LocalTable.CirculatedCount less than Node.CycleCount, set LocalTable.Circulated to false.

Next, Node.NextHopTable is calculated from the receive tables stored in PeerTables.Tables following the procedure specified in L.2.1.10. This information **shall** be passed up to the higher layers of STANAG 5066.

The Receive() procedure returns Token.

L.4.7. Determining Successor

This section sets out the approach for determining the successor, which is used prior to making a general purpose transmission. The detailed approach to determining successor is an implementation choice. This section notes two approaches.

1. Determining Successor by Calculating Best Ring is described in Section L.4.7.1. This is a complex approach, which is expected to give work well for simple and complex topologies.

2. Determining Successor by Link Quality Analysis is described in Section L.4.7.2. This is a simpler approach which will work well for simple networks and is expected to generally work well for target networks.

Either of these algorithms **may** be used or a different algorithm **may** be chosen.

If the last node in TOL.Current is the local node, then TOL.current specifies a ring. If this is the case and the successor is determined to be the penultimate node in TOL.Current, then the ring is unchanged. In this situation TOL.CurrentTransmitted is set to the local node, followed by TOL.Current with the final (local) node removed.

In other situations, TOL.CurrentTransmitted is set to the local node, followed by TOL.Current. Then TOL.CurrentTransmitted is truncated to the last unique node.

L.4.7.1. Determining Successor by Calculating Best Ring

In this approach, all nodes and node connectivity are considered to calculate the best ring. The successor is chosen based on using the best ring.

Care needs to be taken when changing from a working ring. The following should be avoided:
1. Changes to ring where the new ring is not significantly better.
2. Changes to ring with small improvements when the ring has only recently become stable.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

In determining the best ring, preference should be given to an option which reaches the last node in TOL.current in the smallest number of hops.

The key inputs to calculating the best ring transmission order are the receive tables, stored in LocalTable.Table and PeerTables.Tables. L.2.1.10 specifies calculation of an adjacency matrix from this information, which can be used to determine which nodes are directly connected. The adjacency matrix lists bidirectional links, which are essential for token transfer giving recommended maximum transmission speed in each direction. A ring **shall** be determined using bidirectional links.

Calculation of the best ring order is not in general possible; This problem is analogous to the well known travelling salesman problem. A node **may** choose any algorithm when determining "best ring". It is anticipated that good solutions can determined for typical networks in a reasonably straightforward manner. The following baseline algorithms are suggested:

1. For a node with direct connectivity to only one other node, place this node in the ring between two instances of the node to which it connects.
2. If the "outer" node on this list only connects to one other node, place this third node on both ends of the ring being built. This process is repeated to handle nodes connected in a line. Such a line of nodes, with one end connected, is treated as a single node in subsequent calculation.
3. Where a node has only two direct connections, the natural ring fragment is built.
4. If there are multiple fragments, they are joined to form a ring. If end points do not connect directly, L.2.1.10 specifies how to form a route between a pair of nodes. Preference should be given to nodes not yet in the right being built and to fast links.
5. Other nodes can be joined into the ring, seeking first pairs of adjacent nodes to which each remaining node connects to.

There is scope for optimization, to minimize RCL and to use fastest links, noting that links may have different speeds in each direction. It is anticipated that implementation experience will provide input to extend and refine the algorithm specified here, to be updated in a future version of this specification. If there are nodes that cannot be reached from the current node, they are not included in the ring.

L.4.7.2.   *Determining Successor by Link Quality Analysis*

A simpler approach is to consider only TOL.Current and direct links.

The directly connected links can be determined from the receive tables, stored in LocalTable.Table and PeerTables.Tables. L.2.1.10 specifies how to determine a list of bidirectional direct links with transmit quality for each link. Only bidirectional links **shall** be considered.

The following rules are used to select the next link:

1. The predecessor **shall** only be used as the successor when this is the only viable link.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

2. If the local node appears more than once in TOL.current, only nodes prior to the oldest occurrence of the local node **shall** be considered (ignoring the last node if that is the local node).   This rule ensures sensible routing in partially connected networks.

3. For the nodes being considered, use the node with the best transmit quality.

4. If nodes have equal transmit quality, use the node that was least recent in TOL.current.

L.4.8.   Sending Token, WTRP Messages and User Data

A node will transmit data using the Transmit() procedure, which invokes the process described in this section. This is called from the state machine.

Transmission will be made to the node determined by following the procedure of Section L.4.7.

When a node has the token, it will make a transmission that includes the at least one copy of the token and may include user data. The token is encoded as an EOW and will usually be repeated many times in the transmission. Mechanisms to facilitate this are set out in L.6.

A transmission containing WTRP messages **shall** contain the token in each WTRP message. All WTRP messages and tokens in a transmission will have the same source and destination address. Other D_PDUs directed to the successor **may** contain the token.   When token is transferred, it must be put into at least on D_PDU. If necessary, a Padding D_PDU can be used for this. The token will generally be repeated many times.

If Node.CloseRing is true or Node.OperatorDrop is true, set TOL.CurrentTransmitted to empty and do not send any user data in the transmission.

If TOL.CurrentTransmitted is the same as TOL.Transmitted[token.receiver] then the Token Only EOW is used. Otherwise the Token and TOL EOW is used.

token.inviterID is set to Invite.Next.

The following WTRP messages **must** also be transmitted.   These messages may be repeated.

1. If Token and TOL EOW, message.TOL **shall** be sent containing the TOLCurrentTransmitted.   It is recommended to send this message several times.
2. If LocalTable.TableChanged is true, message.TABLE is sent containing LocalTable.Table. LocalTable.Circulated is set to true and LocalTable.CirculatedCount is set to Node.cycleCount.
3. For each received table in PeerTables.ReceivedDirect a message.TABLE is sent.
4. If Node.CloseRing is true, message.CLOSE is sent.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

User data, if available, **may** be sent in this transmission in addition to the messages above which **shall** be sent.

L.4.9. Controlling Invitations

The token indicates which node is next due to make an invitation. This section specifies the algorithm for the node to determine whether or not to issue an invitation and setting the next inviter

These algorithms are specified in the context of two procedures that can be called form the state machine.

ReadyToSendInvite ()

> Return Type: Boolean

> Description: Returns true if it is the local node's turn to invite and if the criteria here are met

If Invite.Next is not the local node, return false.

There is a configurable maximum ring size (MAX_NET_SIZE).   If the number of notes in the ring is equal to or greater than this size, no invitation is issued.   Procedure returns false.

The following parameters are used to control issuing and invitation:

1. Time since Invite.LastTime. If greater than MIN_INVITE_INTERVAL (configurable), an invitation should issued, and procedure returns true.
2. Number of ring circuits since last invite by this node, determined by Node.CycleCount – Invite.LastCycleCount.   If greater than or equal to MIN_INVITE_CYCLES (configurable), invite should be issued and procedure returns true
3. InviteNumber.   If this is less than or equal to EARLY_INVITE_COUNT (configurable) number and ring circuits less than EARLY_INVITE_CYCLES (configurable), an invite should be issued and procedure returns true. This option **may** be omitted. It is designed to give a higher invitation rate in a new ring.

If none of the above conditions are met, no invite should be issued and procedure returns false.

SendInvite()

> Return Type: None

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Description: Send an Invite

Send a message.INVITE with:

- message.INVITE.destination = broadcast address
- message.INVITE.TOL = TOL

Update variables as follows:

- Invite.LastInviteTime set to current time
- Invite.LastCycleCount set to Node.CycleCount
- increment Invite.Number
- Invite.Next set to Node ID of ring member with the value immediately above the local node. If the local node is the highest value, set to the Node ID of ring member with lowest value. This process cycles the ring members, without making any assumptions about ring order.

L.4.10.   Handling Transmissions from Nodes not in the TOL

Where a transmission and WTRP messages are received from a node not in the ring and not simply joining the ring, there are three possible scenarios identified.

1. A node forming a self ring that has not yet heard this ring. Strategy is to just let it find the current ring and join.

2. A node joining elsewhere in the current ring. This will sort out without any action.

3. Another formed ring. The approach is to close the current ring, which will enable a merged ring to form, based on the other ring. Care needs to be taken with rings which are on the edge of communication, because of potential instability due to poor links. This algorithm requires repeat hearing.

The definitive indication of another ring is token transfer. If this is not detected, the other transmission is ignored. If token transfer is detected, this will be recorded. If NUM_OTHER_RING_HEARD (configurable) of token transmissions are heard within OTHER_RING_TIME (configurable), this is considered definitive detection of another ring, which is within range.

When another ring is definitively detected, set Node.CloseRing to true.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.4.11. Overall State Diagram

Figure L-12 below shows the complete state diagram of WTRP. Each state is described in detail in the following sections. The set of states are divided into three subsets that correspond to a station unaffiliated with any token ring, a station soliciting membership in a ring that has invited it to join, and a station in operation within an active token ring.
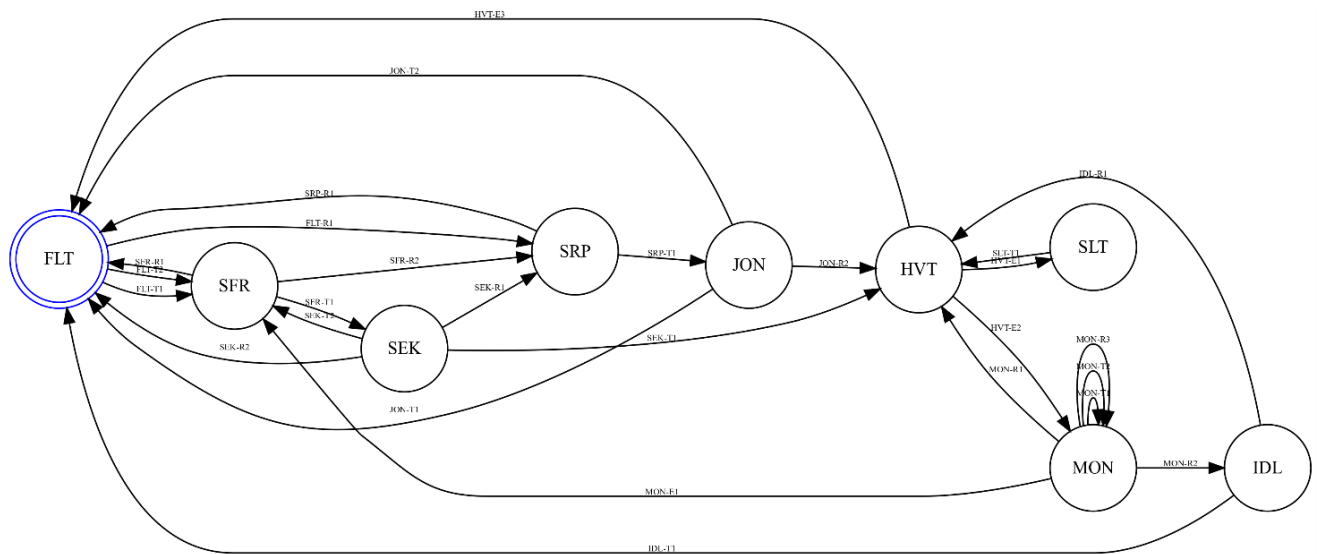
Figure L-12– Overall State Machine

L.4.12. State-Machine Specification

This section and its subsections specify the actions of the WTRP state machine. For every state there are state-entry actions and an outbound-transition table defined. When entering a state, a station first **shall** execute the state-entry actions and then it **shall** wait for an event to occur which triggers one of the transitions to the next state defined in the outbound-transition table. There are two types of event that trigger state transitions:

- Events caused by timeouts; a timeout event is prefixed with the label "Exp" (i.e., for expiry or expired), followed with the name of the timer causing the timeout.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

- Events caused by received data; this event is prefixed with the label "Rcv" (i.e., for 'received'), followed with information to clearly identify what is received.

Only one transition rule **shall** be executed after an event: this **shall** be the first <u>and only the first</u> transition for which the condition is met as the state-machine logic examines the outbound-transitions in the order in which they are listed in the table.

When an action causes data to be transmitted, transition to the next state is expected to happen after the data has been transmitted and before any data is received.

L.4.12.1.   *Floating State (FLT)*

The *Floating State* is the WTRP start-up state and the state in which a station is not part of a ring and waits to join a ring. The floating state is a *listening-only* state. A station **shall** stay in the FLT state until there is a joining opportunity (i.e., a message.INVITE is received inviting the station to solicit membership in the ring) or the TCLT timer expires.

The TCLT timer is used to determine when a station **will** assume that there is no existing ring present. If this timer times out (i.e., expires) the station **shall** proceed to *Self Ring State* (SFR).

If the station receives a message.INVITE it **shall** transit to the *Joining State* (JON) via the *Solicit Reply* (SRP) state,

L.4.12.1.1. Floating State entry actions

On this state entry the station **shall** execute the following actions:

- Execute the IntializeNode() function

- Start the TCLT timer.

- Start Receive()

L.4.12.1.2. Floating State outbound transitions

Outbound transitions from the FLT state are shown in the table below. The most significant outbound transition is to the Solicit-Reply (SRP) State, which occurs when a node receives an invitation to join a ring. For other transitions, the node either remains in the FLT state waiting to receive invitations or transits to the SFR state where it will wait before deciding to send its own invitations for nodes to join its ring.

If any transmission is heard, it will cause the node to wait longer using the configurable TCLT_TRANS_HEARD timer, as there may be an active ring.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Table L-5 - FLT outbound transitions

transition	event	Condition	Action	next state
FLT-R1	EOT	Message.INVITE	Save message.INVITE for processing in SRP state	SRP
FLT-R2	EOT		Start TCLT_TRANS_HEARD Timer Receive()	N/A
FLT-T1	Exp: TCLT			SFR
FTL-T2	Exp: TCLT_TRANS_HEARD			SFR

Figure L-13 - FLT Outbound transitions

L.4.12.2. *Self Ring State (SFR)*

In this state a station has concluded there is no ring to join and therefore will setup a new ring by itself. This condition is called the *self ring*. This state is like the Floating State (FLT) in that it is a listening-only state. On state entry the TSLS timer **shall** be set to a random timeout value, which is used to avoid collisions with any other station trying to setup a ring. If this timer (i.e., if TSLS) times out the station **shall** transit to the *Seeking State* (SEK), sending an invitation to nodes to join its network as an outbound action during the transition.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

While waiting for the TSLS timeout, a message.INVITE **might** be received from another station in SEK state; in this case the station **shall** transit to the *Solicit Reply State* (SRP), where it will respond to the invitation.

If any other transmission is received the station **shall** go to the *Floating State* (FLT).

L.4.12.2.1. Self Ring State (SFR) entry actions

On SFR-state entry a station **shall** start its TSLS Timer with a random time-out value over a configurable range.   Receive() procedure.

L.4.12.2.2. Self Ring State (SFR) outbound transitions

Outbound transitions for the SFR state are shown in the table below. Transitions from the SFR state are triggered: when the station receives an invitation from another node and then will reply; when the station receives any other D_PDU from another station and thus should wait for an invitation; or when the station hears nothing for some time and then sends its own invitation.

Table L-6 - SFR-State Outbound-Transition Table

transition	event	condition	action	next state
SFR-R1	EOT			FLT
SFR-R2	EOT	Message.INVITE	Save message.INVITE to be processed in SRP state	SRP
SFR-T1	Exp: TSLS		SendInvite() See L.4.9	SEK

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Figure L-14 - Outbound transitions from SFR state

L.4.12.3.   Seeking State (SEK)

The *Seeking State* is a state in which a station in a *self ring*, the seeking node, broadcasts a message.INVITE inviting new ring members and listens for replies from solicitors until the TSLW timer times out. The initial TOL for the self ring will have a single entry, which is included in the message.INVITE. Any solicitor will reply by sending a message.JOIN which includes a receive table as message.JOIN.Table. A solicitor may send multiple message.JOIN messages when merging rings. Each message.JOIN message received is processed, which will lead to an updated TOL.

When the TSLS timer expires and no message.JOIN has been received, the node reverts to FLT state. If one or more message.JOIN has been received, the TOL will be updated to reflect all of the new ring members. The station will then proceed to the RCT state.

If a message.INVITE is received before the TLSW timer expires, then another ring (possibly a self-ring) is active within radio range of the station. The station in this case **shall** cease its invitations to other nodes to join its self-ring and **shall** transit to the SRP state, with intent to join the ring it has detected.

Reception by a station when it is in the SEK state of any other transmission shall force the station into the FLT state, as reception of such tokens is an indication there is an active ring within radio range, the self-ring condition no longer applies and the station should not be soliciting to form its own ring.

L.4.12.3.1. Seeking State entry actions

On state exit from the SFR state the station has broadcast a message.INVITE, and shall start the TSLW timer waiting for replies. Start Receive() procedure.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.4.12.3.2. Seeking state outbound transitions

Outbound transitions from the SEK state are shown in the table below. As noted above, reception of message.JOIN does not force an outbound transition; the message.JOIN messages are stored for later processing and the station remains in the SEK state. Reception of a message.INVITE forces a transition to the SRP solicit reply state and reception of any other transmission forces a transition to the FLT floating state; both of these triggering events invalidate the station's assumption that it can form its own ring but with different responses by the station. As long as the station receives only message.JOIN messages, it will continue to handle them until the TLSW timer expires, at which point the station proceeds to the RCT state and subsequent operation in a multi-node network. If no responses of any kind are heard after waiting for replies, the station returns to the SFR self-ring state.

Table L-7 - SEK state outbound transitions

transition	event	condition	action	next state
SEK-R1	EOT	Message.INVITE	InitializeNode() Save message.INVITE for processing in SRP state	SRP
SEK-R2	EOT D_PDUs Received set to true			FLT
SEK-R3	EOT	Message.JOIN	Receive() Note that Message.JOIN is processed within Receive() following L.4.4	N/A
SEK-R4	EOT	D_PDUs Received set to false	Receive()	N/A
SEK-T1	Exp: TSLW timer	At least one message.JOIN received		HVT
SEK-T2	Exp: TSLW timer	No message.JOIN received		SFR

NATO UNCLASSIFIED
Releasable to the Public
**DRAFT FOR REVIEW**

ANNEX L
STANAG 5066
Edition 4

Figure L-15 - Outbound transitions from SEK state

L.4.12.4. *Solicit Reply State (SRP)*

In the *Solicit Reply State* a station will reply to an invitation to join the network (i.e., to a message.INVITE received from another node) and attempt to join the ring.

The TSRP timer **shall** be started on state entry to determine the moment to reply.

While waiting in the *Solicit Reply State* for the TSRP time to expire, other message.JOIN messages might be received and should be ignored, as they could be expected as replies by other stations to the message.INVITE.

Receipt of other message traffic indicates that the channel is not clear during the protocol's invite-and-solicit dialog for new ring members (an error condition in the protocol). The error condition is best handled by having all solicitors return to a known state (the FLT state) in which they do not transmit.

The soliciting station **shall** reply by sending a message.JOIN to the inviting station, the message.INVITE originator. The message.JOIN contains a receive table, which will enable the inviting node to add the local node to the TOL.

L.4.12.4.1. State entry actions:

Set the TSRP timer following the algorithm defined in Section L.5. Start Receive() procedure

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.4.12.4.2. Solicit Reply State Outbound Transitions

Outbound transitions from the solicit transition can be categorized as error recovery or incremental success in joining the network. If any message other than a message.JOIN is heard, the station effectively declares error and transitions to the FLT state, where it will wait for another invitation to join. If the channel remains clear of unexpected traffic, the node sends its solicitation (a message.JOIN) to join the network and transits to the JON state where it will wait to see if its solicitation succeeded.

Table L-8 - SRP outbound transition table

| transition	Event	condition	action	next state
SRP-R1	EOT	Joins Received set to false and D_PDUSs Receives set to true (Message.JOINs being sent by other nodes in response to invite are ignored)		FLT
SRP-R2	EOT	(other than SRP-R1)	Receive()	N/A
SRP-T1	Exp: TSRP		Send message.JOIN where message.JOIN.sender = Node.address message.JOIN.receiver = address of inviter message.JOIN.Table = LocalTable.Table	JON

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Figure L-16 - Outbound transitions from SRP state

L.4.12.5.   *Joining State (JON)*

In the *Joining State* a station has replied to solicitation opportunity by sending a message.JOIN and is now waiting for the *right-to-transmit token*.

The test that the soliciting station has successfully joined the ring is straightforward; the soliciting station either receives an RTT token or receives an updated TOL sent to another station that includes the local node. The TOL is determined by the inviting station and the local (soliciting) station does not know where it will be placed in the TOL.

If the node determines that it is in the TOL, a longer TCON_IN_RING timer is set, as the node can confidently wait for the token to arrive.

L.4.12.5.1. Joining State entry actions

Start the TCON timer.   Start Receive().
.

L.4.12.5.2. Joining State outbound transitions

Outbound transitions for the JON state are shown in the table below. A station either succeeds with its solicitation to join, and transits to the HVT state as a ring member and able to transmit, or fails and transits to the FLT state where it will wait for the next invitation from a ring member (or its own declaration that there is no ring present and eventual transition to the SFR self-ring state).

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Table L-9 - JON-State Outbound-Transition Table

| transition	event	condition	action	next state
JON-R1	EOT	isMember(Node.address) is true	Set TCON_IN_RING timer Stop TCON timer Receive()	N/A
JON-R2	EOT	Token		HVT
JON-R3	EOT	isMember(Node.address) is false	Receive()	N/A
JON-T1	Exp: TCON timer			FLT
JON-T2	Exp: TCON_IN_RING timer			FLT

Figure L-17 - Outbound transitions from JON state

L.4.12.6.   *Have Token State (HVT)*

In the *Have Token State* a station has received the *right to transmit token* (RTT) and as part of the state exit action the channel "owner", is allowed to send D_PDUs as long as the length of its transmission does not exceed the *maximum time to transmit*.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.4.12.6.1. Have Token State entry actions

The node will have received the token in the previous state, which causes the node to enter HVT. The received transmission will have been processed according to L.4.6.

If map or receive table information is needed, but not received, token is sent to predecessor with message.RETRANS and node enters MON state.

The node will verify whether or not an invitation should be sent, following the procedure in L.4.9. If it is determined that an invitation should be sent, and invitation is sent and the node transitions to SLT state. After SLT, the node will return to HVT.

After any invitation has been sent, the token, optional messages and optional user data will be sent following L.4.8, and the node transitions to MON state. L.4.8 will also handle three special situations:

1. If it has been identified that the ring needs to be closed, due to another ring, node will remove itself from TOL and send message.CLOSE to successor along with any messages, but no user data.
2. The HVT state is where the node processes operator or other local requests for the node to drop out of the ring. The node will remove itself from the TOL and send messages but no user data.
3. Where retransmission of Map or TOL is needed.

L.4.12.6.2. Have Token State outbound transitions

Outbound transitions from the HVT state are shown in the table below.

Transitions from the HVT state to the MON state happens after the transmission has completed.

Table L-10 - HVT-State outbound transition table

transaction	event	condition	action	next state
HVT-E1	State entry	ReadyToSendInvite () == true	SendInvite()	SLT
HVT-E2	State entry	ReadyToSendInvite () == false && RingMemberCount() > 1	Transmit() following L.4.8	MON
HVT-E3	State entry	ReadyToSendInvite () == false && RingMemberCount() <= 1		FLT

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Figure L-18 - Outbound transitions from HVT state

L.4.12.7.   *Monitoring State (MON)*

The *Monitoring State* is a state in which a station has finished transmitting data, passed the RTT token to its successor and is waiting for an acknowledgement. An implicit acknowledgement mechanism is used. Any D_PDU received in the MON state with source address equal to the station's successor are taken as implicit acknowledgement that the successor received the RTT token and has taken control of the radio channel. Also, any token transfer on the current ring indicates that the token is circulating in the ring. This second mechanism addresses when the transmission by the successor is not heard.   This monitoring is handled by L.4.5 and reflected in the variable Node.TokenTransferred.

In a two node ring, the token will be passed directly back.

If a TOL has not been received when expected, a message.RETRANS will be received with the token. The token and all messages, but no user data, shall be transmitted following L.4.8.

At expiration of the TPST timer the token transfer is considered to have failed. The station shall resend it for a number of times until the MAX_TOKEN_PASS count has been reached. It is recommended that retransmissions are done at progressively slower speeds.

If the MAX_TOKEN_PASS count has been reached the station declares the next-hop node unreachable and reflects this in connectivity by using the breaklink() procedure.   The Transmit() procedure of L.4.8 is followed again.   This process continues until the token is transferred or there is no possible node to transfer to. If all other nodes are eliminated, the node transitions to SFR.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.4.12.7.1. Monitoring State entry actions

On state entry the station **shall** start the TPST timer, setting the time the station shall wait for receipt of an explicit or implicit acknowledgement of the *right-to-transmit*.

For any entry other than a self-transition, i.e, if the preceding state is not the MON state, Node.txTokenCounter is set to 1.

Start Receive()

L.4.12.7.2. Monitoring State outbound transition table

Outbound transitions from the MON state are shown in the table below.

Transitions representing a successful pass of the *right-to-transmit* token are to HVT state (token comes back) or to IDL state.

Other transitions will cycle through the MON state to transmit to nodes in turn.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Table L-11 - MON-State Outbound-Transition Table

transition	event	Condition	Action / Comments	next state
MON-E1	State entry	RingMemberCount() < 2	If TOL has been reduced to 1, move to self ring	SFR
MON-E2	State entry	Node.NewSucessor == true (set in MON-T2)	Set Node.NewSuccessor to false Transmit token TOL and other messages, and optionally user data following L.4.8	N/A
MON-R1	EOT	Token		HVT
MON-R2	EOT	Node.TokenTransferred == true	Implicit Ack	IDL
MON-R3	EOT	Message.RETRANS	RetransmitMessage(message.RETRANS)	MON
MON-R3	EOT	Node.TokenTransferred == false	Receive()	N/A
MON-T1	Exp: TPST	Node.txTokenCounter < MAX_TOKEN_PASS	Increment Node..txTokenCounter Transmit token TOL and other messages, and optionally user data following L.4.8	MON
MON-T2	Exp: TPST	Node.txTokenCounter >= MAX_TOKEN_PASS	Set Node.txTokenCounter to 1 BreakLink(Node.Successor) OptimizeTOL() Set Node.NewSuccessor to true MON-E1 will check validity of TOL MON-E2 will do the transmission	MON

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Figure L-19- MON Outbound transitions from monitoring state

L.4.12.8. *Idle State (IDL)*

In the *Idle State* the station waits for the *right-to-transmit token* until its TIDL timer expires. When it receives an RTT token, it transits directly to the HVT state. In the IDL state, the node listens for traffic following L.4.4, L.4.5, L.4.6 and L.4.10.

If the TIDL timer expires the station assumes an error-condition in the protocol and transits to the FLT state to recover.

L.4.12.8.1. Idle State entry Actions

The station **shall** start its TIDL timer on state entry. Receive()

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.4.12.8.2. Idle State (IDL) outbound transitions

Outbound transitions from the IDL state are shown in the table below.

Table L-12 - IDL State Outbound-Transition Table

| transition	event	condition	action	next state
IDL-R1	EOT	Token		HVT
IDL-R2	EOT		Receive()	N/A
IDL-T1	Exp: TIDL			FLT

Figure L-20 - Idle State outbound transitions

L.4.12.9.   *Soliciting State (SLT)*

The purpose of the SLT state is to allow new stations to join the ring. A station in the *Soliciting State* is a ring member and has received the *right-to-transmit* from the previous HVT state. A message.INVITE was sent on transition to SLT.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

The station waits until TSLW timer times out and then returns to the HVT state.   Every message.JOIN is processed, leading to update of the TOL, which will be used to determine transmission in HVT state.

L.4.12.9.1. Soliciting State (SLT) entry actions

On state entry the station **shall** start the TSLW timer and listen for message.JOIN.

Start Receive()

L.4.12.9.2. Soliciting State (SLT) outbound transitions

Outbound transitions from the SLT state are shown in the table below.

The normal transition (SLT-T1) to the HVT state occurs when the TSLW time expires, and is made whether or not new stations have solicited to join.

Table L-13 - SLT-State Outbound-Transition Table

transition	event	condition	action	next state
SLT-R1	EOT	Message.JOIN	Receive() Note that Message.JOIN is processed within Receive() following L.4.4	N/A
SLT-R2	EOT		Receive()	N/A
SLT-T1	Exp: TSLW			HVT

**Figure L-21 - Outbound transitions from soliciting state**

L.5.   TIMERS AND PARAMETERS

This section reviews all of the parameters and timers.   The core of this section is a table with all of the specified timers and parameters.   The rationale for the parameter and considerations for optimal settings are described. Each of these parameters and timers **shall** be modifiable by an operator to change the protocol

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

responsiveness or behaviour in response to different operational requirements or tradeoffs (e.g., increased collision probability for newly joining nodes versus reduced solicitation overhead).

L.5.1.   HF Considerations

HF provides a difficult channel with highly variable speed, with short and extended period where no communication is possible. The nature of this channel is a primary consideration in choosing parameter and timer values.

L.5.2.   Surface Wave Considerations

A major target for WTRP is surface wave, particularly for communication between naval vessels where surface communication is good and extends a long way beyond line of sight.   WTRP is less suitable for HF Sky Wave, where it will generally be preferable to use ALE to optimize point to point links, rather than attempting to use a single channel for many nodes. For this reason, default parameters are chosen on the basis of using surface wave to communicate between a group of ships.   Where WTRP is used in other scenarios, it is likely to be desirable to change defaults.

With surface wave, SNR will systematically decrease with distance.   This means that SNR will change as nodes move and will be better for nodes that are closer together.   For this reason, it is anticipated that the variable speed operation provided by this specification will be highly beneficial to good performance.

At the limit of communication distance, communications is expected to be very intermittent.   Care needs to be taken with this, as there is possibility to spend resource trying to maintain a link between a pair of nodes where they can hear each other, but the channel is too poor for robust transfer.

L.5.3.   Ring Forming Strategy

A node that is not in a ring will attempt to join or form a ring.   A node cannot belong to more than one ring, and having two "intersecting" rings is highly undesirable and should be avoided.   If a node hears an existing ring, the strategy is to wait patiently to join the ring.
If a node is not hearing anything, it is hard to distinguish between the following scenarios:
1. No ring, and potentially other nodes that could form a ring.
2. A large ring with long transit time.
3. An existing ring that is hidden by poor HF conditions

L.5.4.   Ring Maintenance Strategy

Once a node is in a ring, efforts should be made to maintain the ring.   Node failure must be allowed for, but is unlikely.   The most common failure is expected to be due to poor HF conditions, and retry is the best approach here.   It is also possible that surface wave distance is increasing between the nodes, leading to the link becoming not being viable

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

L.5.5.   Parameter and Timer Table

Table L-14 sets out parameters and timers.   Some timers specify use of jitter to vary the time to prevent nodes deadlocking with repeat cycle.

Table L-14 – Parameters and Timers

Parameter Name	Default Value	States	Notes
TCLT	2 minutes	FLT	Claim Token Timer (TCLT) - Timer used in the floating-state (FLT). Controls the time a station waits while in the floating state to claim a token before transiting to another state; a station restarts its TCLT timer when it transits to the FLT state. A node in this state does not know if there are other rings. Primary strategy is to find another ring. If this timer is set too short, risk is that two nodes will form a ring while there is another ring. Set too long, and it will delay setting a ring. Making this value slightly longer than an expected typical ring cycle time is sensible.
TCLT_TRANS_HEARD	6 minutes	FLT	This is similar to TCL, except that a transmission has been heard (that is not message.INVITE) and it can be inferred that a ring exists. In this situation, it makes sense to simply wait for an invitation to be issued. This is likely to take a number of ring cycles and a loner timer is appropriate. It is likely that the ring will continue to be heard before an invite is issued, so this timer will generally get reset.
TSLS	n/a	SFR	Solicit Successor Timer (TSLS) - This timer is used in the self-ring-state (SFR). If it times out the station shall transit to the seeking-state (SEK). The timer shall be set with a random timeout to reduce the probability of collisions by transmissions from stations attempting to establish different rings at the same time. This timer is derived from the following algorithm Random Number in range 0 to NUM_SEK_SLOTS multiplied by SEK_SLOT_TIME
NUM_SEK_SLOTS	3	SFR	Number of asynchronous seeking slots; a parameter to calculate the TSLS timer. Collisions are not expected to be common, so a fairly small number is reasonable.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Parameter Name	Default Value	States	Notes
SEK_SLOT_TIME	5 seconds	SFR	Time of the slots for TSLS timer. There is no slot co-ordination, so the slot should be somewhat longer than the expected message.INVITE transmission time
TSRP	n/a	SRP	Solicit Reply Timer (TSRP) - Timer used in the solicit-reply-state (SRP). A station starts its solicit-reply timer when it transits to the SRP state. If it expires a station will transit to the JON state where it replies to one of the received SLS tokens, if any, with a SET token. This timer is randomized, as it is possible that multiple nodes will hear message.INVITE and send message.JOIN in response This timer is derived from the following algorithm Random Number in range 0 to NUM_SLS_SLOTS multiplied by SEK_INTERVAL
SLS_INTERVAL	5 seconds	SRP	Time of the slots for TSRP timer. There is no slot co-ordination, so the slot should be somewhat longer than the expected message.INVITE transmission time
NUM_SLS_SLOTS	5	SRP	Number of asynchronous seeking slots; a parameter to calculate the TSRP timer. Collisions are not expected to be common, so a fairly small number is reasonable.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Parameter Name	Default Value	States	Notes
TCON	n/a	JON	Contention Timer (TCON) - Timer used in the joining-state (JON). It controls the time a station waits for a response from another station following an attempt to join the network, so-named because failure to receive a response is attributed to contention with other stations attempting to join the network at the same time; a station restarts its contention timer when it goes to JON state. This needs to be tied to the TSRP timer, as the time to wait is linked to the time which the inviting node will wait before it sends. Then the length of a transmission needs to be waited. The transmission will either advertise the TOL (enabling detection of ring joing) or will pass the token to the local node. The following algorithm is used: (NUM_SLS_SLOTS +1) * SLS_INTERVAL + MAX_TX_TIME This gives time for all nodes to send responses. It assumes that all nodes on the network have same values for the two parameters used here.
TCON_IN_RING	n/a	JON	The node has determined that node is in ring and needs to wait for its turn. This allows the maximum time around the entire ring, with a factor that allows for other nodes being added and retransmissions. MAX_TX_TIME * RCL() * TCON_FACTOR
TCON_FACTOR	1.2	JON	Factor to adjust TCON_IN_RING timer for retransmissions and TOL changes.
TIDL	n/a	IDL	Idle Timer (TIDL) - Timer used in the idle-state (IDL). It controls the time a station waits for its right-to-transmit before transiting to the floating-state (FLT). The following algorithm is based on max time for token to propagate around the ring. MAX_TX_TIME * RCL() * TIDL_FACTOR
TIDL_FACTOR	1.2	JON	Factor to adjust TIDL timer for retransmissions and TOL changes.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Parameter Name	Default Value	States	Notes
TSLW	n/a	SLT, SEK	Solicit Wait Timer (TSLW) - This timer is used in the soliciting state (SLT) and the seeking-state (SEK) by stations inviting new ring members. When it expires, the inviting station it will update the Transmit Order List. The following algorithm is used, which reflects the algorithm used is SRP state to transmit message.INVITE: (NUM_SLS_SLOTS +1) * SLS_INTERVAL
TPST	3 seconds	MON	Token Pass Timer (TPST) - Used in the monitoring-state (MON). It controls the time a station waits after passing an RTT (or other) token to another station and failing to hear an implicit acknowledgement before considering the right-to-transmit as lost. It is expected that the this will happen very quickly in normal conditions. If this fails, either the token was not received or the transmission from next node is not being heard. For the former condition a short value is best. Operational experience is expected to give insight on the best choice for this setting and others.
MAX_TOKEN_PASS	3	MON	Specifies the number retransmissions of a *right-totransmit* after which a station shall stop trying the current node and update the TOL to give a different node.
MAX_NET_SIZE	255	HVT	The maximum number of nodes allowed in the ring. If there is a fixed number of nodes in operation, setting this value will save the overhead of issuing invitations when the ring is "complete". Setting a low limit will improve ring performance, but is unfair to excluded nodes. The recommended default is to use the upper bound, which in practice means no limit.
MIN_INVITE_INTERVAL	15 minutes	HVT	Minimum interval between which a node will issues invites. Setting this value high will improve performance. Setting it low will enable new nodes to join more quickly.
MIN_INVITE_CYCLES	4	HVT	Minimum number of complete stable ring cycles between a node issuing invites. . Setting this value high will improve performance. Setting it low will enable new nodes to join more quickly.
EARLY_INVITE_COUNT	2	HVT	If a node has issued this number of invites or less, use EARLY_INVITE_CYCLES

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Parameter Name	Default Value	States	Notes
EARLY_INVITE_CYCLES	1	HVT	Minimum number of complete stable ring cycles between a node issuing invites when EARLY_INVITE_COUNT is used
RECEIVE_TABLE_EXPIRY_AGE	n/a	HVT	This sets the time when entries should be removed from the receive table. A node is expected to transmit at least once per ring cycle. This timer is tied to the maximum time for a ring cycle and a factor, which allows for a number of cycles and that cycle time may be much shorter than the maximum. The following algorithm is used MAX_TX_TIME * RCL() * RTEA_FACTOR
RTEA_FACTOR	3	HVT	Variable to control RECEIVE_TABLE_EXPIRY_AGE
MAX_TX_TIME	127.5 seconds	HVT, MON	Maximum Transmit Time. This is a general STANAG 5066 control parameter that can be up to 127.5 seconds. By limiting this time, ring transit times will be reduced and latency reduced. However, for bulk transfers, longer times will significantly improve throughput, particularly where only a small proportion of the nodes are sending data.
NUM_OTHER_RING_HEARD	3	Active Ring	Number of transmissions that must be heard from another ring before closing current ring. Setting this too high will delay ring merging and may lead to two nodes transmitting together. Setting it too low may lead to ring tear down when it is not viable to establish a merged ring.
OTHER_RING_TIME	10 minutes	Active Ring	Period within which NUM_OTHER_RING_HEARD transmissions from another ring must be heard before closing current ring. Considerations for setting this parameter are similar.

L.6.   TOKEN AND MESSAGE TRANSMISSION

L.6.1.   Transmitting Tokens

Tokens are carried in EOWs. Every D_PDU carries exactly one EOW. A token can be carried in any D_PDU which is being sent to the successor, which will include all WTRP messages being sent. STANAG 5066 may require to use EOWs for other purposes. Subject to this, it is desirable to send multiple copies of the token for reliability.

Additional EOWs may be sent using the ACK_PDU or the PADDING D_PDU specified in STANAG 5066 Annex C.   This can enable extra token copies to be sent, which may be particularly useful if most D_PDUs being sent are to different destinations.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX L
STANAG 5066
Edition 4

Where possible, the Tokens should be spread out over the length of the transmission, to minimize the effects of fading and data loss.

L.6.2. Transmitting Messages

L.4.8 specifies which messages are required to be sent. Messages are small and it may be beneficial to repeat them, particularly if a higher transmission speed is chosen. Where Map and Table messages are critical, it is particularly desirable to duplicate.

Where messages are duplicated, it is best to have them at different parts of the transmission, rather than close together.

All messages are processed after the transmission is received, so there is no particular benefit to placing them in a particular part of the transmission.

L.6.3. Speed and Interleaver Selection

WTRP can be operated at fixed speed or at variable speed. This specification provides a recommended maximum transmission speed for each node. When making a transmission, the speed associated with the "slowest" node needs to be considered as maximum speed for the whole transmission. Where speed is highly variable, this may impact the choice of which D_PDUs to send.

When there is a lot of data to send, it is recommended to use the maximum transmission speed with a long or very long interleaver. This is expected to maximize throughput and reliability.

Messages will often be WTRP messages only, which are generally small. Where there is a smaller amount of data a slower speed and shorter interleaver is likely to give better performance.

When sending message.INVITE to a broadcast address, it is recommended to use a conservatively slow speed. Similarly with message.JOIN. Once communication is established, both ends can use SNR to determine an appropriate operational speed.

L.7. CHANGES IN EDITION 4

The overall model and service provided by Annex L is unchanged. The state machine has some small changes. The protocol in Edition 4 is completely different to Edition 3. This change was made because of significant problems with the protocol in Edition 3.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

ANNEX N –
GUIDANCE ON ADDRESS MANAGEMENT IN
STANAG 5066 NETWORKS
(INFORMATIVE)
DRAFT FOR REVIEW: Version 1.0 of 20th April 2020

N.1   INTRODUCTION

This annex provides guidance for allocation and use of STANAG 5066 addresses in operational systems. It defines a block addressing scheme [1] that can be used for formally planned and ad-hoc operations with a minimal amount of system reconfiguration   required.

N.2   MANAGED ADDRESSES

This annex **does not** alter the requirements for size, format, or representation of node addresses defined elsewhere in this STANAG, notably but not limited to the   following:

- Annex A– Node ADDRESS Encoding for all   [S_]Primitives;

- Annex C– Size of Address [encoding in   DPDUs];

- Annex C – Source and Destination Address [encoding in   DPDUs];

Addresses in this annex are given in the dotted-decimal format defined [2] in Annex   C, and abbreviated generically in the form w.x.y.z

N.3   NODE ADDRESS-BLOCK ASSIGNMENTS

Only full-length (i.e., 28-bit) S'5066 node addresses **shall** be subject   to   the assignment recommendations of this annex.

Regional blocks **shall** be defined by the 4 most-significant bits of the full-length address (i.e., the first element w of the dotted-decimal form w.x.y.z).

---

[1] This Annex bases its address-block allocation strategy on a proposal put forth by the US Navy as part of their concept of operations for the Battle-Force E-Mail 66 system, which is one of the larger deployments of STANAG-5066 based systems managed under a single operational authority.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

Node addresses less than full-length (i..e, with w = 0) **may** be assigned and managed on an ad hoc basis in unique situations; they are otherwise unmanaged by any provision of this Annex. Management authorities for less-than-full-length addresses are not   defined.

Managed regional address blocks **shall** be defined by the 4 most-significant bits of   the full-length address, in accordance with the Table   below.

Table N-1 — Top-Level Address-Block Assignments by Region

Address Range	Regional Assignee	Management POC	Comment
0.0.0.0 — 0.255.255.255	unassigned	ad-hoc	Variable-length addresses are unmanaged
1.0.0.0 — 1.255.255.255	United States	US DoD	includes US Armed Forces and Homeland Security as major S'5066 users
2.0.0.0 — 2.255.255.255	North America	TBD	
3.0.0.0 — 3.255.255.255	North America	TBD	other than US government
4.0.0.0 — 4.255.255.255	South America	TBD	
5.0.0.0 — 5.255.255.255	NATO	TBD	
6.0.0.0 — 6.255.255.255	Europe	TBD	
7.0.0.0 — 7.255.255.255	Europe	TBD	
8.0.0.0 — 8.255.255.255	Asia	TBD	
9.0.0.0 — 9.255.255.255	Asia	TBD	
10.0.0.0 — 10.255.255.255	Africa	TBD	
11.0.0.0 — 11.255.255.255	Middle East	TBD	
12.0.0.0 — 12.255.255.255	Australasia, New Zealand, and Oceana	TBD	
13.0.0.0 — 13.255.255.255	Non-Governmental Organizations	TBD	includes, e.g., the International Committee of the Red Cross
14.0.0.0 — 14.255.255.255	other	TBD	
15.0.0.0 — 15.255.255.255	other	TBD	

Authority to assign and manage a specific address for a specific system **may** be devolved (perhaps more than once) to a regional or sub-regional administrator. Devolution of management authority within a regional block is beyond the scope of this   Annex.

A system's address need not — indeed, should not — be changed when it moves   from one geographic region to another. Regional membership and responsibility to assign an address to a system is based on the region in which the system is registered or affiliated (both of these processes are outside of the scope of the STANAG), rather than the system's location. The major exceptions to this provision are systems belonging to nations involved   in NATO operations, where the option of using a national or NATO address may be exercised. Even in this case, however, the election to use an address in one block or another is based on administrative or operational criteria, not

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

physical location.

Whether or not an address manager is assigned for a given administrative region, nodes **should** be assigned addresses from a regional block with which they are registered or affiliated. This will reduce the risk of assigning an address already in use, even absent more active administrative co-ordination.

Entities or individuals wishing to nominate themselves for the role of regional / sub- regional address administrator should contact the NATO C3 Staff (Attn: Network Domain Branch), NATO Headquarters, Brussels, BE.

Within the global-regional blocks, specific 5066 addresses **shall** be assigned to individual nations using the second element "x" of the dotted-decimal address form w.x.y.z. Multiple values of "x" can be pre-assigned in order to give Nx65535 address nodes per nation or organizational unit. Exceptions to this general rule are the US Government and NATO, which will likely be large-scale implementers of the STANAG and are assigned their own unique regional "w" blocks.

Sub-regional blocks allocations are defined in the subsections below. The lists of nations is based on the "Independent States in the World" (URL: http://www.state.gov/www/regions/independent_states.html) as released by the US Office of the Geographer and Global Issues, January 21, 2000.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

N.3.1   **North-American National Address Schema**

Sub-regional assignments for North American national and organizational entities are defined below.

The US Government, as a large scale implementer and deployer of S'5066 systems,   is allocated the 1.x.y.z address block and organizational address assignments within this block **shall** conform to the table below.

Table N-2 — North American (US Government) National Addressing Schema   (1.x.y.z)

"1" . "x"	Organization	"1" . "x"	Organization
1.1	US Navy	1.8	US Joint Forces Commands
1.2	US Marine Corps	1.9 – 19	Other US Military
1.3	US Air Force	1.20	US Federal Emergency Management Agency
1.4	US Army	1.21	BATF
1.5	US Special Operations Forces	1.22	US Federal Bureau of Investigation
1.6	US Coast Guard	1.23 – 1.255	Other US Govt
1.7	US Military Sealift Command		

S'5066 system address allocations for North American nations ― excluding systems registered with the US Government ―**shall** conform to the table below.

Table N-3 — North American (non-US Government) National Addressing Schema   (2-3.x.y.z)

"w" . "x"	Country	"w" . "x"	Country
2.1	Antigua & Barbuda	2.15	Guatemala
2.2	Bahamas	2.16	Haiti
2.3	Barbados	2.17	Honduras
2.4	Belize	2.18	Jamaica
2.5, 2.6, 2.7	Canada	2.19	Mexico
2.8	Costa Rica	2.20	Nicaragua
2.9	Cuba	2.21	Panama
2.10	Dominica	2.22	St. Kitts and Nevis
2.11	Dominican Republic	2.23	St. Lucia
2.12	Ecuador	2.24	Trinidad & Tobago
2.13	El Salvador	3.1	United States (non-US Government)
2.14	Grenada		

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

N.3.2   **South American National Address Schema**

S'5066 system address allocations for South American nations **shall** conform to the table below

Table N-4 — South American National Addressing Schema (4.x.y.z)

"w" . "x"	Country	"w" . "x"	Country
4.1	Argentina	4.7	Guyana
4.2	Bolivia	4.8	Paraguay
4.3	Brazil	4.9	Peru
4.4	Chile	4.10	Suriname
4.5	Colombia	4.11	Uruguay
4.6	Ecuador	4.12	Venezuela

N.3.3   **NATO Address Schema**

NATO is expected to be a large user of S'5066 systems and consequently has its own top-level address block assigned to it. Sub-allocations within this address block are made to NATO nations and commands and **shall** conform to the table below.

Table N-5 — NATO Addressing Schema (5.x.y.z)

"w" . "x"	Country	"w" . "x"	Country
5.1	Belgium	5.12	Norway
5.2	Czech Republic	5.13	Poland
5.3	Denmark	5.14	Portugal
5.4	France	5.15	Spain
5.5	Germany	5.16	Turkey
5.6	Greece	5.17	United Kingdom
5.7	Hungary	5.18	United States
5.8	Iceland	5.19	Allied Command for Operations
5.9	Italy	5.20	Allied Command for Transformation
5.10	Luxembourg	5.21	NATO CIS Services Agency
5.11	Netherlands	5.22	NATO Consultation, Command and Control Agency

Allocations within the NATO block are made for national forces provided in support of NATO operations, e.g., to a NATO Response Force (NRF), Combined Joint Task Force (CJTF) or Standing NATO Force (SNF). National forces participating in a NATO mission or operation **should** use addresses in the NATO block. However, nations belonging to NATO **may** elect to   use addresses within their own allocated block instead the NATO allocation. The option to use a nationally administered address or a NATO administered address for any given system

should be resolved as part of the operational planning for any mission in which the system participates.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

N.3.4 **European National Address Schema**

S'5066 system address allocations for European nations **shall** conform to the table below.

Table N-6 — European National Addressing Schema (6-7.x.y.z)

"w" . "x"	Country	"w" . "x"	Country
6.1	Albania	6.27	Lithuania
6.2	Andorra	6.28	Luxembourg
6.3	Austria	6.29	FYROM[3]
6.4	Belarus	6.30	Malta
6.5	Belgium	6.31	Moldova
6.6	Bosnia & Herzegovina	6.32	Monaco
6.7	Bulgaria	6.33	Netherlands
6.8	Croatia	6.34	Norway
6.9	Cyprus	6.35	Poland
6.10	Czech Republic	6.36	Portugal
6.11	Denmark	6.37	Romania
6.12	Estonia	6.38, 39, 40	Russia
6.13	Finland	6.41	San Marino
6.14, 6.15	France	6.42	Slovakia
6.16, 6.17	Germany	6.43	Slovenia
6.18	Greece	6.44	Serbia & Montenegro
6.19	Holy See	6.45	Spain
6.20	Hungary	6.46	Sweden
6.21	Iceland	6.47	Switzerland
6.22	Ireland	6.48	Turkey
6.23, 6.24	Italy	6.49	Ukraine
6.25	Latvia	6.50, 51, 52	United Kingdom
6.26	Liechtenstein	6.53	Yugoslavia

---

[3] Turkey recognises the Republic of Macedonia by its constitutional name

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

N.3.5 Asian National Address Schema

S'5066 system address allocations for Asian nations **shall** conform to the table below.

Table N-7 — Asian National Addressing Schema (8-9.x.y.z)

"w" . "x"	Country	"w" . "x"	Country
8.1	Bangladesh	8.12	Malaysia
8.2	Bhutan	8.13	Maldives
8.3	Brunei	8.14	Mongolia
8.4	Cambodia	8.15	Myanmar
8.5	China	8.16	Nepal
8.6	India	8.17	Phillipines
8.7	Indonesia	8.18	Singapore
8.8	Japan	8.19	Sri Lanka
8.9	Korea, North	8.20	Taiwan
8.10	Korea, South	8.21	Thailand
8.11	Laos	8.22	Vietnam

N.3.6 African National Address Schema

S'5066 system address allocations for African nations **shall** conform to the table below.

Table N-8 — African National Addressing Schema (10.x.y.z)

"w" . "x"	Country	"w" . "x"	Country
10.1	Algeria	10.27	Lesotho
10.2	Angola	10.28	Liberia
10.3	Benin	10.29	Madagascar
10.4	Botswana	10.30	Malawi
10.5	Burkina Faso	10.31	Mali
10.6	Burundi	10.32	Mauritania
10.7	Camaroon	10.33	Mauritius
10.8	Cape Verde Islands	10.34	Morocco
10.9	Central African Republic	10.35	Mozambique
10.10	Chad	10.36	Namibia
10.11	Comoros	10.37	Niger
10.12	Congo (Brazzaville)	10.38	Nigeria
10.13	Congo (Kinshasha)	10.39	Rwanda
10.14	Côte d'Ivoire	10.40	Sao Tome
10.15	Djbouti	10.41	Senegal
10.16	Egypt	10.42	Sierra Leone
10.17	Equatorial Guinea	10.43	Somalia
10.18	Eritrea	10.44	South Africa
10.19	Ethiopia	10.45	Sudan

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

10.20	Gabon		10.46	Swaziland
10.21	Gambia		10.47	Tanzania
10.22	Gambia		10.48	Togo
10.23	Ghana		10.49	Tunisia
10.24	Guinea		10.50	Uganda
10.25	Guinea-Bissau		10.51	Zambia
10.26	Kenya		10.52	Zimbabwe

N.3.7 Middle East National Address Schema

S'5066 system address allocations for Middle Eastern nations **shall** conform to the table below.

Table N-9 — Middle East National Addressing Schema (11.x.y.z)

"w" . "x"	Country		"w" . "x"	Country
11.1	Afganistan		11.12	Lebanon
11.2	Armenia		11.13	Oman
11.3	Azerbiajan		11.14	Pakistan
11.4	Bahrain		11.15	Qatar
11.5	Georgia		11.16	Saudi Arabia
11.6	Iran		11.17	Syria
11.7	Iraq		11.18	Tajikistan
11.8	Jordan		11.19	Turkmenistan
11.9	Kazakhstan		11.20	United Arab Emirates
11.10	Kuwait		11.21	Uzbekistan
11.11	Kyrgyzstan		11.22	Yemen

N.3.8 Australia, New Zealand, and Oceania National Address Schema

S'5066 system address allocations for Australia, New Zealand and nations in Oceania **shall** conform to the table below.

Table N-10 — Australia, New Zealand and Oceana National Addressing Schema   (12.x.y.z)

"w" . "x"	Country		"w" . "x"	Country
12.1	Australia		12.9	Papua New Guinea
12.2	Fiji		12.10	Samoa
12.3	Kiribati		12.11	Seychelles
12.4	Marshall Islands		12.12	Solomon Islands
12.5	Micronesia		12.13	Tonga
12.6	Nauru		12.14	Tuvalu
12.7	New Zealand		12.15	Vanuatu
12.8	Palau			

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX N
STANAG 5066
Edition 4

N.3.9 Non-Governmental Organization (NGO) and Other Address Schema

S'5066 system address allocations for Non-Governmental Organizations and other entities **shall** conform to the table below.

Table N-11 — Non-Governmental Organization and Other Addressing Schema   (13-15.x.y.z)

"w" . "x"	NGO		"w" . "x"	NGO
13.1	United Nations			
13.2	Red Cross			
14.0-255	Other			
15.0.255	Other			

N.4 Changes in Edition 4

Minor updates to align text to rest of STANAG

Removal of Ether-Client address management, as no longer applicable.

Annex O: HF Operator Chat
(Optional)

DRAFT 1.0 of 14th April 2020

HF Operator Chat (HFCHAT) provides a simple and efficient chat communication between operators for exchange of short IA5 text messages. It can also be used to broadcast message to all operators. If provided, the HFCHAT client **shall** conform to the minimal requirements defined herein. The HFCHAT client is intended as a simple mechanism to allow subnetwork operators to test and coordinate their system configurations. It is not a general purpose chat system.

O.1.   General Requirements for HFCHAT

HFCHAT clients **shall** use the ITA5 / ASCII character set to exchange short text messages between subnetwork operators.

HFCHAT messages **shall** consist of sequences of characters terminated by a carriage-return/line-feed pair (i.e., terminated by the octet-pair 0x0D, 0x0A).

The HFCHAT messages length, i.e., the number of octets in the character sequence and including terminating carriage-return/line-feed pair, **shall** not exceed the subnetwork MTU size.

In general, methods of presentation and display to the operator of HFCHAT messages, as well as methods for text-message entry, are beyond the scope of this STANAG and left as implementation options. The following implementation guidelines are recommended, however:

- HFCHAT clients **should** provide a common entry and display area for HFCHAT messages.

- HFCHAT clients **should** provide a short, viewable history of previous messages sent and received (e.g., of the last N messages, N a value (configurable or not) in the range [10,100]).

- HFCHAT clients **should** provide an indication of the source of any HFCHAT messages that it receives (e.g., by displaying the STANAG 5066 address of the originator of the HFCHAT message)

- HFCHAT clients **should** provide an indication of the time-of-receipt of any HFCHAT messages that it receives.

- HFCHAT clients **should** provide confirmation-of-delivery (node delivery) indications for HFCHAT messages it sends when they are sent using ARQ delivery service.

- HFCHAT clients **should** provide both 1:1 and broadcast transmission.

O.2.   Subnetwork Service Requirements

An HFCHAT client **shall** bind to the HF Subnetwork at SAP ID 5.

Priority of messages **should** be configurable and **may** be operator-selectable. Default priority **shall** be 12.

HFCHAT clients **may** use either point-to-point or point-to-multi-point addressing modes to send HFCHAT

messages.

The default subnetwork-service requirements when using the point-to-point addressing mode **shall** be as follows:

- Transmission Mode =   ARQ
- Delivery Confirmation =   NODE DELIVERY
- Deliver in Order =   IN-ORDER DELIVERY

For point-to-point operation, the address in the primitive will be an individual node address corresponding appropriately to the HF subnetwork address of the remote HFCHAT client to which the message is addressed.

The default subnetwork-service requirements when using the point-to-multipoint addressing mode **shall** be as follows:

- Transmission Mode =   non-ARQ; number of repeats (configurable)
- Delivery Confirmation =   NONE
- Deliver in Order =   NOT IN-ORDER DELIVERY

For point-to-multipoint operation, the address in the primitive will be a multicast-address corresponding to the group of remote HFCHAT clients to which the message is addressed. Establishment of the multicast address for the group may be done through standard operating procedure or out-of-band channel. Option to use the default broadcast address **should** be provided.

The message's sequence of characters **shall**[1] be bit- and byte-aligned with the octets in an S_Primitive's U_PDU, with the least-significant bit (LSB) of each character aligned with the LSB of the octet. The unused eighth (i.e, MSB) of the octet **shall**[2] be set to zero.

HFCHAT messages **shall** be encapsulated and sent within S_PRIMITIVES, one message per S_PRIMITIVE, with the message-terminating carriage-return/line-feed pair encapsulated in the S_PRIMITIVE as the last two octets of the U_PDU field.

O.3.   Changes Since Edition 3

Minor clarifications and corrections made.

Annex P: ACP 127 & Character-Oriented Serial Stream

(Optional)

DRAFT 1.0 of 14th April 2020

The Character-Oriented Serial Stream (COSS) is primarily intended for operation of ACP 127 formal messaging over STANAG 5066 ARQ service. COSS is specified as a generic service, so that it cold be used by other applications with similar transport requirements. The ACP 127 protocol transmits a stream of ITA2 or IA5 characters over a serial line. The COSS service provides a service equivalent to serial line transmission. It provides a replacement for ACP 127 operation directly over an HF modem.

P.1. CHARACTER-ORIENTED SERIAL STREAM (COSS) CLIENT

This annex defines a character-oriented serial-transport service for the HF subnetwork. This provides a reliable emulation of a serial line type service.

The character-oriented serial stream (COSS) service **may** be used in place of other HF serial transport services, for example a simple modem. To provide high reliability and end-to-end assurance of data delivery, the COSS client uses the STANAG 5066 ARQ mode to provide higher data reliability than simple transmission over a conventional modem might afford.

COSS defines two modes of operation:

1. Integrated Operation

2. Baseband Operation

An implementation of this annex **may** offer one or both modes.

P.2. Integrated Operation

COSS may be integrated directly with ACP 127 or another application as shown in Figure P-1.
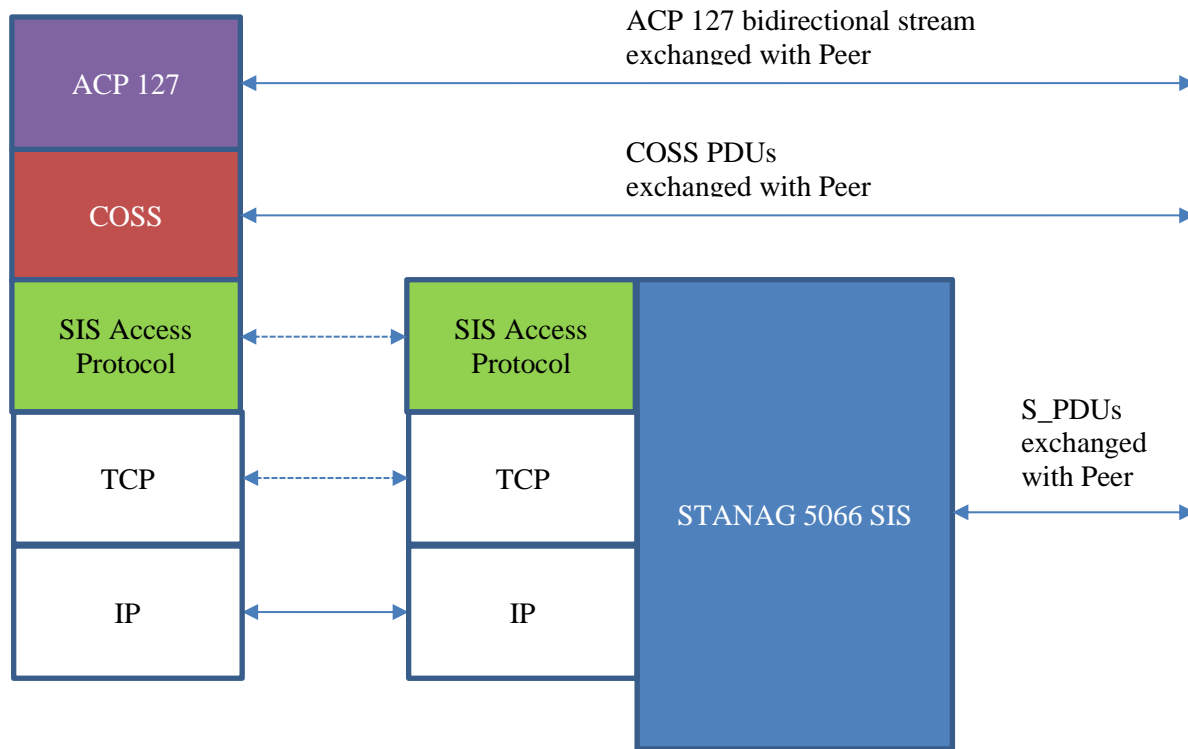
Figure P-1: ACP 127 Integrated Mapping onto COSS

Figure P-1 shows how an application can use a COSS client directly to communicate with a STANAG 5066 server using the SIS Access Protocol. The service provided by the COSS layer enables peer ACP 127 servers to exchange a stream of data in each direction.

P.3. Baseband Operation

COSS can also be configured with a serial interface, so that the COSS module interfaces to ACP 127 or other applications using a serial interface and communicates with a STANAG 5066 server using SIS Access Protocol as specified in Annex S. This mode is termed "baseband operation", as it can directly replace an application using a baseband serial connection directly to an HF modem.

The interfaces for the Character-Oriented Serial Stream (COSS) Client using Baseband operation are as shown in the Figure below.

Figure P-2: COSS Client Interfaces

Requirements for the COSS client are placed on its baseband serial interface, its interface to the HF Subnetwork Interface Sublayer (SIS), and its internal processing. Implementations of the COSS client following Baseband operation **shall** use a baseband serial interface, and implementations **shall**[2] be in accordance with the requirements stated herein.

P.3.1. Base-band Serial Interface

A baseband mode COSS client **shall** provide a baseband serial interface meeting the following requirements.

a. *Physical/Electrical*: one of following physical interfaces **shall** be provided:

1. Signalling and connector conforming to EIA/RS-232, EIA//RS-530, configured as Data Communications Equipment (DCE);

2. Signalling and connector conforming to V.35 DCE.

b. *Serial Transmission Mode*: all of the following transmission modes **shall** be supported (as configuration options):

1. Asynchronous: 1 Start Character, selectable 5, 6, 7, or 8 data bits, and 1 or 2 stop bits.

2. Synchronous: provide/accept clock at 1x Data Rate (other rate-multipliers **may** be supported);
HDLC line control protocol (other synchronous line protocols **may** be supported in addition to HDLC)

c. *Flow-Control*: all of the following flow-control disciplines **shall** be supported (as configuration options):

1. RTS/CTS, DTR/DTS hardware handshaking;

2. XON/XOFF software flow-control;

3.   None.

P.4. Character Sets Supported by COSS

The character sets shown in the Table below **shall** be supported by a COSS client:

Table P-1: Character Sets Supported by COSS

Character Set	Comments
ITA-2 / Baudot	5-bit character sets: w/ asynchronous protocol = ITA-2; w/ synchronous protocol = Baudot
6-bit codes	6-bit character codes
ITA-5 / ASCII	~ASCII - 7-bit character set [NB: the 7.0-unit 64-ary ITA-2 code defined in STANAG 5030 is encoded in this form, in accordance with section F.3.4.4.1]
Octet data	Any data presented in arbitrary 8-bit formats

P.5. Subnetwork Service Requirements for COSS

COSS clients **shall** bind to the HF Subnetwork at SAP ID 1.

A COSS client **shall** submit its PDUs to the HF subnetwork using the S_UNIDATA_REQUEST Primitives defined in Annex A of this STANAG.

The address in the primitive **shall** be a STANAG 5066 address corresponding to the HF subnetwork address of the host at which the destination COSS client(s) is/are located.

The default service requirements defined when the client binds to the subnetwork **shall** be as follows:

1.   Transmission Mode =   ARQ
2.   Delivery Confirmation =   NONE
3.   Deliver in Order =   IN-ORDER DELIVERY

P.6. Data Encapsulation Requirements

The characters from the character stream **shall** be encapsulated within S_PRIMITIVEs using any one of the modes described herein. Implementation of all modes is **mandatory** for a COSS client.

Selection of any given mode for operation **shall** be a configuration parameter in a COSS client, and dependent on the character-set for which the COSS client is configured.

Selection of any given mode **must** be coordinated at the sending and receiving node for use on a given link, through either standard operating procedure or out-of-band coordination channel.

P.6.1. Encapsulation of Arbitrary Octet Data

Octet data in any arbitrary format for the COSS client **shall** be byte-aligned with the octets in each U_PDU encapsulated in the S_UNIDATA_PRIMITIVE.

The least-significant bit (LSB) of each character received on the serial interface **shall** be aligned with the LSB of the octet.

P.6.2. Encapsulation of ITA-5

Characters in ITA-5 format (or other 7-bit character format such as ASCII) for the COSS client **shall** be aligned with the octets in each U_PDU encapsulated in the S_Primitive, one character per octet, as follows.

	MSB			Octet Data				LSB
Octet / Bit	7	6	5	4	3	2	1	0
Octet$_I$	0	msb	-	-	$ITA5_j$	-	-	lsb

The least-significant bit (LSB) of each 7-bit character received on the serial interface **shall** be aligned with the LSB of the octet.

The value of the MSB bit of each octet **shall** be set to zero for ITA5 Encapsulation.

P.6.3. Encapsulation of ITA-2

Two methods of encapsulation of ITA-2 characters are defined:

1. 'Loose-Pack ITA2 Encapsulation' (LPI2E), and

2. 'Dense- Pack ITA2 Encapsulation' (DPI2E).

A COSS client **shall** implement both methods of ITA2 encapsulation.

Selection of either method **must** be coordinated by sending and receiving node for use on a given link, through either standard operating procedure or out-of-band coordination channel.

P.6.4. 'Loose-Pack Encapsulation of ITA-2 characters

The Loose-Pack' ITA-2 Encapsulation (LPI2E) algorithm **may** be used for any character set represented as 5-bit symbols. It transports one 5-bit symbol in each octet within the U_PDU field of S_primitives, using the basic packing arrangement defined in the Figure below.

	MSB			Octet Data				LSB
Byte / Bit	7	6	5	4	3	2	1	0
Octet$_I$	0	0	0	msb	-	$ITA2_j$	-	lsb

Figure P-3: Loose-Pack Encapsulation of ITA-2 Characters

The least-significant bit (LSB) of each 5-bit character received on the serial interface **shall** be aligned with the LSB of the octet.

The value of the three most-significant bits of each octet **shall** be set to zero for LPI2E.

P.6.5. 'Dense-Pack' Encapsulation of ITA-2 characters

The 'Dense-Pack' ITA-2 Encapsulation (DPI2E) algorithm **may** be used for any character set represented as 5-bit symbols. It efficiently transports three 5-bit symbols in a pair of octets, called an Encapsulation Pair, within the U_PDU field of S_primitives, using the basic packing arrangement defined in the Figure below.

	MSB			Octet Data				LSB
Octet / Bit	7	6	5	4	3	2	1	0
i = 2n	$ITA2_{j+1}$	-	lsb	msb	-	$ITA2_j$	-	lsb
i = 2n +1	DP_FLG	msb	$ITA2_{j+2}$ (defined only if DP_FLG = 1)			lsb	msb	$ITA2_{j+1}$

Figure P-4: Nominal Dense-Pack ITA-2 Encapsulation Pair

P.6.5.1. Encapsulation Pairs

The first octet of an Encapsulation Pair **shall** be an even-numbered octet in an S_Primitive's U_PDU field (i.e., i = {0,2,4,6, …} , with i = 0 the first octet in the U_PDU).

The second octet of an Encapsulation Pair **shall** be an odd-numbered octet in an S_Primitive's U_PDU field (i.e., i = {1,3,5,7, …} , with i = 0 the first octet in the U_PDU).

The MSB of the second octet of an Encapsulation Pair **shall**[1] be the DP_FLG ('dense-pack flag') that indicates whether the Encapsulation Pair contains three ITA-2 characters (DP_FLG = 1) or two ITA-2 characters (DP_FLG = 0). Encoding of these cases **shall**[2] be performed as follows:

- The first case is defined to be a "Three-into-Two Encapsulation Pair", which **shall** be encoded in accordance with this figure

	MSB			Octet Data				LSB
Byte / Bit	7	6	5	4	3	2	1	0
i = 2n	$ITA2_{j+1}$	-	lsb	msb	-	$ITA2_j$	-	lsb
i = 2n +1	DP_FLG = 1	msb	$ITA2_{j+2}$		lsb		msb	$ITA2_{j+1}$

Figure P-5: Three-Into-Two Encapsulation Pair

- the second case is defined to be a "Two-into-Two Encapsulation Pair", which **shall** be encoded in accordance with this figure:

Byte / Bit	MSB			Octet Data				LSB
	7	6	5	4	3	2	1	0

i = 2n	$ITA2_{j+1}$	-	lsb	msb	-	$ITA2_j$	-	lsb
i = 2n +1	DP_FLG = 0	0	0	0	0	0	msb	$ITA2_{j+1}$

Figure P-6: Two-Into-Two Encapsulation Pair

P.6.5.2. **Maintaining Character-Count Integrity with DPI2E**

The 'Dense-Pack' ITA-2 Encapsulation (DPI2E) algorithm **shall not** add or delete ITA-2 characters from the character stream being transported. This property is denoted character-count integrity.

To maintain character-count integrity, the DPI2E algorithm **shall** depend on the number of characters that remain after all initial characters have been encapsulated as 3-into-2 Encapsulation Pairs.

For a buffer of size B characters encapsulated within a U_PDU field of length L, the DPI2E packing algorithm is defined as follows:

1. **Case R = (B modulo 3) = 0.** In this case, no characters remain after all initial characters are packed in Three-into-Two Encapsulation Pairs:

   - all ITA-2 characters **shall**[1] be densely packed as Three-into-Two Encapsulation Pairs. [i.e.,.: Each Encapsulation Pair will contain three ITA-2 characters, with DP_FLG =1].
   - The U_PDU length **shall**[2] be set to the value L = (2 * B/3).

2. **Case R = (B modulo 3) = 1.** In this case, one character remains after all initial characters are densely packed in Three-into-Two Encapsulation Pairs:

   - the first (B-1) ITA-2 characters **shall**[1] be densely packed as 3-into-2 Encapsulation Pairs. [i.e.: each of the Encapsulation Pair will contain three ITA-2 characters, with DP_FLG =1].
   - The last remaining ITA-2 character **shall**[2] be loosely packed in the last octet of the U_PDU in accordance with the 'Loose-Pack Encapsulation of ITA-2 characters specification' of section F.3.4.3.1.
   - The U_PDU length **shall**[3] be set to the value L = (2 * ((B-1)/3) + 1).

3. **Case R = (B modulo 3) = 2.** In this case, two characters remain after all initial characters are densely packed in Three-into-Two Encapsulation Pairs:

   - the first (B-2) ITA-2 characters **shall**[1] be densely packed as 3-into-2 Encapsulation Pairs. [I.E.: Each Encapsulation Pair will contain three ITA-2 characters, with DP_FLG =1]
   - the two remaining ITA-2 characters **shall**[2] be packed as a 2-into-2 Encapsulation Pair. [I.E.: Each Encapsulation Pair will contain two ITA-2 characters, with DP_FLG =0; the bit positions corresponding to the third ITA-2 character in the Encapsulation Pair will be set to zero.].
   - the U_PDU length **shall**[3] be set to the value L = (2 * ((B-2)/3) + 2).

ANNEX P
STANAG 5066
Edition 4

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

P.6.5.3. **Character Unpacking Requirements for DPI2E**

In accordance with standard operation procedure or out-of-band coordination circuit, the receiver **must** be configured to perform Dense-Pack ITA-2 Encapsulation.

The receiving client **shall** perform the inverse DPI2E algorithm to unpack the ITA2 characters the U_PDUs contained within an S_UNIDATA_INDICATION primitive:

For a U_PDU of length L,

- If L = 1, the receiving client **shall** unpack the single ITA2 character from the loosely packed octet and send it to the output serial stream;

- If L > 1 and L even, the receiving client **shall**[1] unpack the ITA2 characters in order from each Encapsulation Pair of octets, continuing in order for each successive Encapsulation Pair in the U_PDU. Unpacked ITA2 characters **shall**[2] be sent in the order in which they are unpacked to the output serial stream. Receiving nodes **may** log a processing error if any Encapsulation Pair except the last has DP_FLG = 0.

- If L > 1 and L odd, the receiving client **shall**[1] unpack the ITA2 characters in order from each Encapsulation Pair of octets, continuing in order for each successive Encapsulation Pair in the U_PDU, and unpacking the last ITA2 character from the single octet (last, and not part of an Encapsulation Pair) in the U_PDU. Unpacked ITA2 characters **shall**[2] be sent in the order in which they are unpacked to the output serial stream. Receiving nodes **may** log a processing error locally if any Encapsulation Pair has DP_FLG = 0. Receiving nodes **may** log a processing error locally if the three most-significant bits of the octet are nonzero.

P.6.6. **Encapsulation of 6-bit Character Codes**

Characters in 6-bit formats for the COSS client **shall** be aligned with the octets in each U_PDU encapsulated in the S_Primitive, one character per octet.

	MSB				Octet Data			LSB
Octet / Bit	7	6	5	4	3	2	1	0

$Octet_I$	0	0	msb	-	$6\_bitChar_j$	-	-	lsb

The least-significant bit (LSB) of each 6-bit character received on the serial interface **shall** be aligned with the LSB of the octet.

The value of the MSB bit of each octet **shall** be set to zero for 6-bit character encapsulation.

The special case of the 64-ary ITA-2 character code defined for STANAG 5030 is specified below.

P.6.6.1. **Encapsulation of 64-ary ITA-2 (i.e., STANAG 5030) character codes**

The 64-ary ITA-2 code as defined in STANAG 5030 for Single and Multicahannel VLF/LF Broadcasts actually uses a 7-bit character. It consists of 6-bits for information plus a 7^{th} bit (the Stop Bit) that does not change. This allows the possibility that some applications may choose a simple approach to shortening the STANAG 5030 64-ary ITA-2 code to a true six-bit code (i.e., by deleting the stop bit from the code, relying on other measures to provide character synchonization).

In 7-bit format (i.e, unshortened), the 7.0 unit STANAG 5030 64-ary ITA-2 code **shall** be encapsulated as specified in section F.3.4.2.

As a shortened 6-bit code (i.e, with the 7^{th}/stop-bit removed, the STANAG 5030 64-ary ITA-2 code **shall** be encapsulated as specified in section F.3.4.4.

As the overhead from both approaches is equivalent, there is no reason with respect to STANAG 5066 operation to shorten the code. Consequently, use of the unshortened 7.0 unit STANAG 5030 64-ary ITA-2 code is preferred. Other external considerations may apply however that would favor use of the shortened code.

P.6.7. **Character-Flush Requirements**

It is assumed that the COSS client will be implemented with an input buffer for temporary storage of characters from the stream prior to their encapsulation in S_PRIMTIVES for transmission over the subnetwork. The events that trigger transfer of characters from this buffer to an S_PRIMITIVE (i.e, that triggers a 'character- flush' operation) need to be specified for the client. Of concern are the performance tradeoffs that exist for various character-flush disciplines. For instance, frequent character-flush operations will reduce end-to-end latency and increase the overhead, while for infrequent character-flush operations, triggered only when the size of the input buffer is the subnetwork's Maximum Transmission Unit Size (MTU), the opposite is true. As this is a largely performance issue and not an interoperability issue, a number of different behaviours could be defined.

The COSS client **shall** have a capability to configure the behaviour of its input-buffer character-flush discipline.

Characters **shall** be flushed from the COSS input buffer and encapsulated in an S_PRIMITIVE for transmission over the subnetwork when one or a combination of the following events occur:

1. The number of characters in the input buffer exceeds a configurable threshold value, COSS_BUF_FLUSH_THRESHOLD [NB: if the specified threshold value is greater than the subnetwork MTU size, then COSS_BUF_FLUSH_THRESHOLD shall equal the MTU value.];

2. A carriage-return/line-feed input character-pair is detected in the input character stream. [NB: this behaviour provides line-by-line transmission of a character stream organised as lines of text.]

3. A configurable timeout interval has occurred following the arrival in the input buffer of the last received character. [NB: this behaviour ensures that characters in the input buffer are eventually transmitted if one of the first two events has not occurred.]

Other behaviours for the character-flush disciplines **may** be defined as additional and configurable implementation options, e.g., triggering a character-flush operation on detection of a user-specifiable character sequence defined as an End-Of-Message (EOM) sequence. [NB: such operation would be useful to support legacy

systems such as the ACP-127 messaging systems, which use a defined character sequence as a message delimiter.]

P.7. Changes Since Edition 3

Edition 4 contains a number of minor corrections and clarifications, but is broadly unchanged.

Direct operation of COSS from an ACP 127 implementation without use of serial line is clarified.

Non-ARQ operation is removed, as this will not work.

Annex Q: ACP 142 (Optional)

DRAFT 1.0 of 15th April 2020

This annex specifies operation of ACP 142 ("P_Mul – A Protocol for Reliable Multicast Messaging in Constrained Bandwidth and Delayed Acknowledgement (EMCON) Environments") over STANAG 5066. This is important for two services.

1. Formal Military Messaging following STANAG 4406, using Tactical Military Messaging protocol specified in STANAG 4406 Annex E.
2. Email using RFC 8494 ("Multicast Email (MULE) over Allied Communications Publication (ACP) 142")

This annex also includes RCOP (Reliable Connection Oriented Protocol) and UDOP (Unreliable Datagram Oriented Protocol) which are used in support of ACP 142 over STANAG 5066.

Q.1.   Overall Architecture
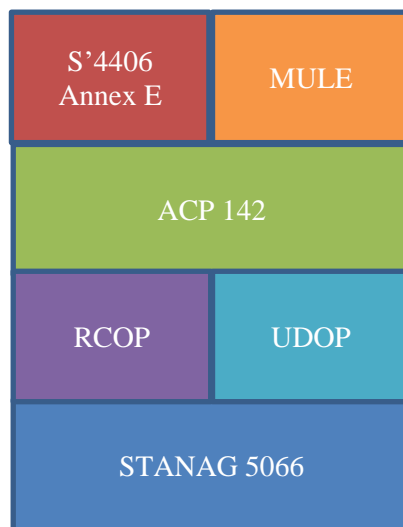
Figure Q-1: ACP 142 over STANAG 5066 Architecture

Figure Q-1 shows the overall architecture for ACP 142 operation over STANAG 5066. ACP 142 has a choice of protocols to interface to STANAG 5066:

1. Reliable Connection-Oriented Protocol (RCOP) is specified in Section Q.5

2. Unreliable Datagram-Oriented Protocol (UDOP) is specified in Section Q.6

ACP 142 is a general purpose multicast transport and can support a range of applications. This annex is focussed on two applications:

1. Formal Military Messaging using STANAG 4406 Annex E over ACP 142. This is described in Section Q.2.

2. Email using RFC 8494 ("Multicast Email (MULE) over Allied Communications Publication (ACP) 142"). This is described in Section Q.3.

Q.2.   STANAG 4406 TACTICAL MILITARY MESSAGE HANDLING

A Formal Military Message is different from an interpersonal message in that it is a message sent on behalf of an organization, in the name of that organization, that establishes a legal commitment on the part of that organization under military law, and has been released in accordance with the policies of the originating nation. Examples are military orders. Individuals may send organizational Messages to other individuals on behalf of their respective organizations. Formal Military Messages support the additional services associated with ACP 127, and is therefore interoperable with ACP 127 systems.

Formal Military Messages are handled by Military Message Handling Systems (also called High Grade Messaging Service). An MMHS takes responsibility for the delivery, formal audit, archiving, numbering, release, emission and distribution of received formal messages on behalf of the originating organization. An MMHS is accountable under military law to provide a reliable, survivable and secure messaging service on behalf of the originating organization. The MMHS fulfills the military messaging service requirements of ACP-121 and offer high standards of messaging reliability and security. The formal messaging service is seen as the vehicle for secure mission critical, operational military applications.

NATO support for the X.400 protocols in Military Message Handling Systems (MMHS), as defined by NATO STANAG 4406, is mandated in the NATO C3 Technical Architecture (NC3TA), Volume 4 - NATO Common Standards Profile (NCSP). STANAG 4406 Annex E (S'4406E) specifies an adaptation of the X.400/X.500 protocols in STANAG 4406 for Tactical Military Messaging Handling Systems (T-MMHS), with cross-references to STANAG 5066 as one example of a low-bandwidth bearer service to which it has interfaces. This annex summarises and specifies additional requirements for the interface between S'4406E-based T- MMHS and a STANAG 5066-compliant HF subnetwork.

There are several classes of Tactical Messaging Interfaces (i.e., peer-to-peer interfaces between messaging systems) defined by S'4406E, but only one of these (TMI-1) is supported by this annex. TMI-1 is the Tactical Messaging Interface between two Light Message Transfer Agents (LMTA) as specified in S'4406E.

The S'4406E protocol stack defines a compact file format that is passed to ACP 142 for multicast transfer to the recipient peer MTAs. This format and the interface to ACP 142 is fully specified in S'4406E.

Q.3.   Multicast Email (MULE)

An equivalent service for standard email is provided by using RFC 8494 ("Multicast Email (MULE) over Allied Communications Publication (ACP) 142"). This specification takes a standard SMTP message and generates a

compressed file for multicast transfer by ACP 142. Compression uses the same format as S'4406E. This is fully specified in RFC 8494.

Q.4. ACP 142 Layer

ACP 142 ("P_Mul – A Protocol for Reliable Multicast Messaging in Constrained Bandwidth and Delayed Acknowledgement (EMCON) Environments") provides a reliable multicast transfer service designed to support messaging. It is used by S'4406E and MULE, with service mappings fully defined in those specifications.

ACP 142 is specified to use a datagram service. By default it uses UDP (User Datagram Protocol) specified in RFC 768. The ACP 142 specification clearly notes that any datagram service can be used.

STANAG 4406 Annex E references mapping using Wireless Datagram Protocol (WDP) defined as part of the Wireless Access Protocol (WAP) specification. WDP defines operation over a variety of transports. As the WAP specification does not include STANAG 5066, the datagram mapping to STANAG 5066 is summarized in STANAG 4406 Annex E Appendix A. This annex specifies additional information on the mapping. If there are any inconsistencies, this annex defines the normative specification.

ACP 142 is mapped onto RCOP and/or UDOP as defined in the following sections. Each ACP 142 PDU is encoded as the User Data of an RCOP or UDOP datagram. RCOP and UDOP traffic **may** be mixed.

Q.4.1. Mapping ACP 142 onto UDOP

ACP 142 datagram services make use of the Unreliable Datagram-Oriented Protocol (UDOP) specified in Section Q.6. UDOP **shall** be used for all multicast communication. UDOP **shall** be used for unicast nodes in EMCON state. UDOP **may** be used for unicast nodes not in EMCON state, but this is **not recommended**.

UDOP is based on the STANAG 5066 non-ARQ service, which allows a configurable number of retransmissions. It is **recommended** to not retransmit and to choose transmission speeds where frame loss is anticipated to be low. Data loss is most efficiently handled by ACP 142 retransmission mechanisms and ACP 142 FEC.

It is **recommended** that the ACP 142 MTU size is chosen so that every ACP 142 PDU is transmitted in a single STANAG 5066 UNIDATA. This gives a maximum ACP 142 PDU size of three bytes less than the STANAG 5066 MTU size (typically 2048 bytes). It **may** be desirable to choose a smaller ACP 142 PDU size in order to optimize performance.

Q.4.2. Mapping ACP 142 onto RCOP

ACP 142 datagram services **may** make use of the Reliable Connection-Oriented Protocol (RCOP) specified in Section Q.5. RCOP defines a datagram service based on the STANAG 5066 ARQ service. It is **recommended** that RCOP is used for unicast non-ARQ transmission because:

1. STANAG 5066 ARQ retransmission is more efficient, as this is done at D_PDU level, where D_PDU size is generally matched to current HF characteristics. ACP 142 retransmission is a the ACP 142 PDU level; and

2. STANAG 5066 ARQ retransmission is faster, because it is part of the protocol. ACP 142 retransmission is based on a number of timers, which need to be set conservatively to avoid data duplication.

Node Delivery confirmation **shall** be selected. Order **shall** be AS_THEY_ARRIVE.

It is **recommended** that the ACP 142 MTU size is chosen so that every ACP 142 PDU is transmitted in an integral number of maximum size STANAG 5066 UNIDATAs. This will lead to optimal use of STANAG 5066.   Use of larger messages will reduce ACP 142 overhead.

 RCOP and the underlying STANAG 5066 services generally provide reliable transfer. However, there are error conditions where this does not happen, and the ACP 142 layer **shall** allow for this loss by use of the ACP 142 timers.   The operational characteristics over UDOP/Non-ARQ and RCOP/ARQ are very different, so it is **recommended** that timers are independently configurable for each of these modes of operation.

Q.4.3.   Addressing

The ACP 142 layer uses its own (four byte) addressing.   These addresses need to be mapped onto STANAG 5066 Addresses, as set out in this section.

Each Unicast ACP 142 address **shall** be mapped onto an equivalent STANAG 5066 node address.

The global broadcast ACP 142 address **shall** be mapped onto the default STANAG 5066 broadcast address.

It is **recommended** that ACP 142 dynamic multicast addresses are **not** used with STANAG 5066. If they are used they shall be mapped onto the default STANAG 5066 broadcast address.

For fully interconnected STANAG 5066 networks, there is no benefit to defining static multicast addresses other than a global broadcast address. However, for partially connected WTRP networks, use of static multicast addresses can lead to performance improvements.   If this is done ACP 142 static multicast addresses **shall** be mapped onto equivalent STANAG 5066 group addresses.

Q.4.4.   Priority Mapping

Delivering data according to priority is important in HF systems, particularly for military use. Messages have a priority, which **shall** be used to set priority in the lower layers. Both STANAG 5066 and ACP 142 layers have priorities. ACP 142 priority is 0-255, with 0 as highest priority.   ACP 142 can be used for both formal military messaging and for email, so priority mappings are defined for both services.   The priority mapping for formal military messaging is defined in Table 1.

Message Priority	ACP 142 Priority	STANAG 5066 Priority
OVERRIDE	3	13
FLASH	5	11
IMMEDIATE	7	9
PRIORITY	9	7
ROUTINE	11	5
DEFERRED	13	3

Table 1: Formal Military Messaging Priority Mapping

Email is often operated without priority, but some email services operate with a three-level priority. The three level mapping is defined in Table 2. For messages without priority defined, the "normal" mapping **shall** be used.

Message Priority	ACP 142 Priority	STANAG 5066 Priority
Urgent	10	6
Normal	12	4
Non-Urgent	14	2

Table 2: Email Priority Mapping

Q.4.5. SAP Selection

SAP assignment is set out in Annex F. There are two ACP 142 assignments, which shall be used for these classes of ACP 142 traffic:

- For Military Formal Messaging using STANAG 4406 Annex E or MULE, SAP 2 **shall** be used.

- For Email using MULE, SAP 7 **shall** be used.

Q.5. RELIABLE CONNECTION-ORIENTED PROTOCOL

This section specifies a simple Reliable Connection-Oriented Protocol (RCOP) for reliable data connections between applications using the ARQ services of the HF subnetwork. RCOP provides a minimal header to support multiplexed connections between a pair of nodes through a single hard or soft link. Applications using an RCOP connection are identified uniquely by a field in the header of the RCOP PDU.

Q.5.1.   RCOP Protocol Data Unit (RCOP_PDU)

The format of all RCOP U_PDUs **shall** be as shown in the Figure below.
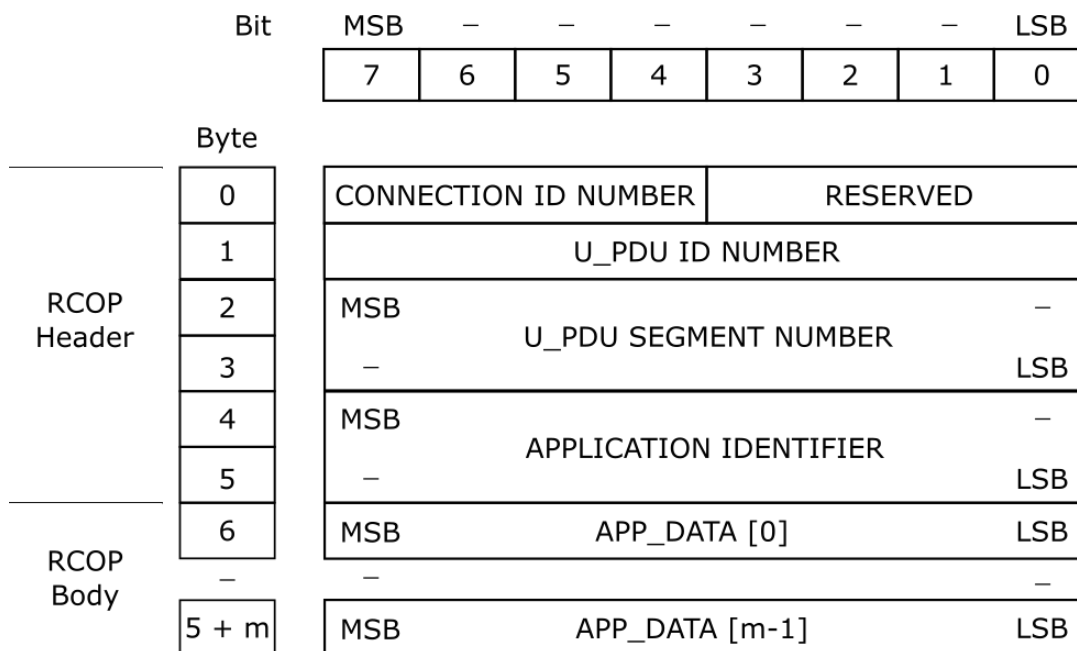

Figure Q-2: Format for Reliable Connection-Oriented Protocol Data Units

The following are required for RCOP PDUs:

1.   RCOP clients may provide multiplexed transport service for more than one application simultaneously by establishing multiple connections, each identified by its CONNECTION_ID_NUMBER field. The CONNECTION_ID_NUMBER field **shall** be a value from 0-15.

2.   New values for the CONNECTION_ID_NUMBER field **shall** be dynamically assigned as new connections are established. Further details regarding assignment and co-ordination of connection ID numbers are not specified here.

3.   Connection ID number 0 **shall** be reserved for non-multiplexed connections or the first multiplexed connection.

4.   The reserved bits **shall** be set to 0.

5.   U_PDU ID numbers **shall** be assigned consecutively to user PDUs (U_PDUs) serviced by the connection.

6.   The U_PDU segment number **shall**[1] be assigned consecutively to segments within a single U_PDU. The first segment transmitted **shall**[2] be assigned segment number 0. If a U_PDU is not segmented, the single segment that is transmitted **shall**[3] be assigned number 0.

7.   The APPLICATION_IDENTIFIER **shall** be assigned in accordance with the requirements of section F.10 of this Annex. This field serves to identify the application (i.e., higher-level protocol) using the connection. End-to-end interoperability can be achieved only if this is a unique value for any given end-user application, otherwise there is the (likely) chance that the applications connected by the RCOP client and subnetwork will be incompatible.

8.   The APP_DATA[] field **shall** contain the m-bytes of application data sent over the connection. Segmentation and reassembly rules for mapping the application's data into the APP_DATA[] field are in general outside of the scope of this STANAG, and are application dependent.

Q.5.2.   RCOP Subnetwork Service Requirements

RCOP clients **shall** bind to the HF Subnetwork at SAP ID 6. Client priority is a configuration- dependent and implementation-dependent parameter.

Each RCOP_PDU sent over the subnetwork **shall** be embedded in an S_UNIDATA_REQUEST primitive, each byte of the RCOP_PDU corresponding to a byte in the U_PDU field of the S_Primitive (see Figure). RCOP_PDUs will be encapsulated in like manner when delivered by the subnetwork to the destination client in an S_UNIDATA_INDICATION.

	APP_DATA[0]	-- APP_DATA --	APP_DATA[m-1]
Bytes[0, 5] (header)	- Byte[6]	- RCOP_PDU - -- - (body)	Byte[5+m]
S_UNIDATA_REQUEST_HEADER	Byte[0]	- - - U_PDU - - -	Byte[k-1]

Figure Q-3: Mapping of RCOP_PDU into an S_UNIDATA_REQUEST Primitive for transmission.

The encoded data for the RCOP client **shall** be bit-/byte-aligned with the octets in each U_PDU encapsulated in the S_UNIDATA_PRIMITIVE, with the least-significant bit (LSB) of each character aligned with the LSB of the octet.

Note that, in accordance with the provisions of Annex A of this STANAG, if the subnetwork interface sublayer receives a S_UNIDATA_REQUEST primitive with an RCOP Protocol Data Unit larger than the maximum MTU size, the S_UNIDATA_REQUEST will be rejected.

An RCOP client **shall** set the default service requirements for S_UNIDATA_REQUEST primitives as follows:

-    Transmission Mode =   ARQ

-    Delivery Confirmation =   NODE DELIVERY or CLIENT DELIVERY

-    Deliver in Order =   IN-ORDER DELIVERY or AS_THEY_ARRIVE

The address in the primitive will be an individual node address corresponding to the HF subnetwork address of the host at which the destination RCOP client is located.

Q.5.3.   RCOP Segmentation and Reassembly Requirements

. Note that, in accordance with the provisions of Annex A of this STANAG, if the subnetwork interface sublayer receives a S_UNIDATA_REQUEST primitive with an RCOP Protocol Data Unit larger than the maximum MTU size, the S_UNIDATA_REQUEST will be rejected. RCOP clients **must** segment their data and place it into the APP_DATA[] field of the PDU accordingly.

If IN_ORDER delivery-order is specified, an RCOP client **shall** simply take the U_PDU segments in the sequence in which they arrive to construct a copy of the original U_PDU sent by the source. If AS_THEY_ARRIVE delivery-order is specified, an RCOP client **shall** be responsible for reassembling the U_PDU segments it receives in proper sequence order.

RCOP clients **may** devise any algorithm of their own choice for segmentation and reassembly, but the RCOP_PDU fields for CONNECTION_ID_NUMBER, U_PDU_ID NUMBER, U_PDU SEGMENT NUMBER
are available for such use. Segmentation and reassembly algorithms **shall** not use the APPLICATION IDENTIFIER field.

Note that, in general, a local RCOP client could be receiving data on two different connections, each established by another remote RCOP client as the remote client's sole connection. In this case, the remote clients would each have specified a connection ID number of zero. As an alternate but similar scenario, two RCOP clients on the same remote node but attached to different SAP IDs could connect to the local RCOP client. In this case also, data could be received with the same connection ID number. Other scenarios in which the same connection ID has been assigned to data received by an RCOP client might also occur, as dynamic connections are made and broken with different nodes. Thus, to reassemble segmented application data without ambiguity, an RCOP client **must** distinguish U_PDU segments for its receive connections using the unique combination of (SOURCE_ADDRESS, SOURCE_SAP_ID, CONNECTION_ID_NUMBER). The SOURCE_ADDRESS is the
address of the originator of the RCOP_PDU, and the SOURCE_SAPID is the SAP_ID to which the remote RCOP client is attached. All three parameters the can be obtained unambiguously from the S_UNIDATA_INDICATION primitive in which the RCOP_PDU is delivered to the client, as can the CONNECTION_ID_NUMBER, U_PDU_ID NUMBER, U_PDU SEGMENT NUMBER.

Q.6.   UNRELIABLE DATAGRAM-ORIENTED PROTOCOL (UDOP)

This section defines a simple Unreliable Datagram-Oriented Protocol (UDOP) using the non-ARQ services of the HF subnetwork, with a minimal header to support multiplexed datagram delivery. Since non-ARQ services are used, the UDOP may support a multicast service through the use of group addresses within the HF Subnetwork. Applications using a UDOP non-ARQ delivery service are identified uniquely by a field in the

header of the UDOP PDU.

Q.6.1. UDOP Data Unit

UDOP Protocol Data Units **shall** be defined and used identically to those defined for the Reliable Connection-Oriented Protocol in Section Q.5.1.

Q.6.2. UDOP Subnetwork Service Requirements

UDOP clients **shall** bind to the HF Subnetwork at SAP ID 7. Client priority is configuration- dependent and implementation-dependent parameters for this application.

A UDOP client **shall** submit its U_PDUs to the HF subnetwork using the normal S_UNIDATA_REQUEST Primitives, with the default service requirements defined as follows:

- Transmission Mode = non-ARQ
- Delivery Confirmation = none
- Deliver in Order = AS_THEY_ARRIVE

The address in the primitive will be an individual node address corresponding to the HF subnetwork address of the host on which the destination message stores are located. The UDOP client at the destination **shall** be responsible for reassembling the U_PDU segments in proper sequence order.

The encoded data for the UDOP client **shall** be bit-/byte-aligned with the octets in each U_PDU encapsulated in the S_UNIDATA_PRIMITIVE, with the least-significant bit (LSB) of each character aligned with the LSB of the octet.

Q.7. Changes in Edition 4

ACP 142 is moved from Annex F to this Annex.

RCOP and UDOP are included in this annex, as they are used by ACP 142 and use by other protocols is not anticipated.

There are no protocol changes to the protocols specified here.

The Annex is formulated as ACP 142, rather than STANAG 4406 Annex E. Technically, this does not lead to a change, but it means that the annex now encompasses other protocols using ACP 142, in particular MULE (RFC 8494).

Only the TM-1 (MTA to MTA) option of STANAG 4406 Annex E is included. Other options are not included, as they are not being used.

Priority mappings for the layers are specified.

Annex R: Routing Sublayer
(Optional)

DRAFT 1.0 of 6th April 2020

The Routing Sublayer is an optional sublayer above Channel Access Sublayer (CAS) and below Subnet Interface Sublayer (SIS). The Routing Sublayer supports transfers of data that needs to traverse one subnet twice or needs to traverse multiple subnets. It also supports systems with multiple independent subnets. The primary goal of Routing Sublayer is to support operation over Wireless Ring Token Protocol specified in STANAG 5066 Annex L when there is partial connectivity.

R.1.   Routing Sublayer Overview

The STANAG 5066 service defines operation over a single subnet. This specification adds a routing layer so that the STANAG 5066 SIS service can be provided across multiple subnets. This is needed to provide data communication between all nodes when Wireless Ring Token Protocol is used with partially connected topology. It also provides additional capabilities for other architectures.

R.2. Goals of the Routing Sublayer

R.2.1. Wireless Token Ring Protocol Support

Wireless Token Ring Protocol (WRTP) as specified in STANAG 5066 Annex L, defines a framework for supporting communication between nodes sharing a channel where the nodes cannot communicate directly. It specifies a MAC level service that provides information on node connectivity to the higher layers of STANAG 5066.

The routing sublayer defined in this specification makes use of this connectivity information to enable users of the STANAG 5066 service to communicate data to nodes that are not directly connected.

R.2.2. Routing across Multiple Subnets

STANAG 5066 service defines access to a single channel (typically access by a single modem) and service to nodes connected by that channel. A system may be connected to multiple channels with independent STANAG 5066 services. This routing sublayer enables transparent interconnection of such services. Examples as to where this might be useful.

1. A ship with two HF connections: Skywave reach back to shore; WTRP surface wave communication to other ships. This routing protocol would enable them to interconnect.
2. A shore station with skywave access to several mobile units, each with a dedicated channel and STANAG 5066 service. The routing sublayer would enable STANAG 5066 communication between the mobile units.
3. A MANET (Mobile Ad hoc Network), with a mix of HF links and line of sight links (e.g., UHF).

R.2.3. Transparent Client Access to multiple Subnets

Where a system has multiple STANAG 5066 subnets, an additional capability that this layer gives is to enable a SIS client to address a STANAG 5066 node without needing to know which STANAG 5066 server is used to access the desired node.

R.3. Alternate Approaches

There are two "obvious" alternatives to the routing sublayer defined here, which are discussed below. Both are highly desirable approaches for some deployments, but they cannot address all scenarios.

R.3.1. Application Relay

When transferring application data over HF, it is highly desirable to proceed on hop at a time. This enables communication to be tuned for one network, without introducing relay. This can be done, and is recommended, whenever routing configuration is stable.
However, where routing can change, for example with WTRP, application relay is not viable.
This approach is useful for some applications such as messaging and XMPP, but other applications do not usefully support application relay.

R.3.2. IP Relay

Use of IP is a common approach used to build MANETs over multiple subnetworks. In many cases, this is a good and flexible choice. However, HF subnetworks have high and variable latency, which leads to problems deploying these applications over HF. Use of relay over multiple networks would compound these problems. This IP approach will not be generally viable for HF, particularly "bulk" applications using TCP. Analysis is provided in the Isode White Paper <u>Measuring and Analysing STANAG 5066 F.12 IP Client</u>.

R.4. Routing Layer Architecture

This section provides a recap of the STANAG 5066 Ed3 architecture, and shows how the Routing Sublayer (RS) extends this architecture.

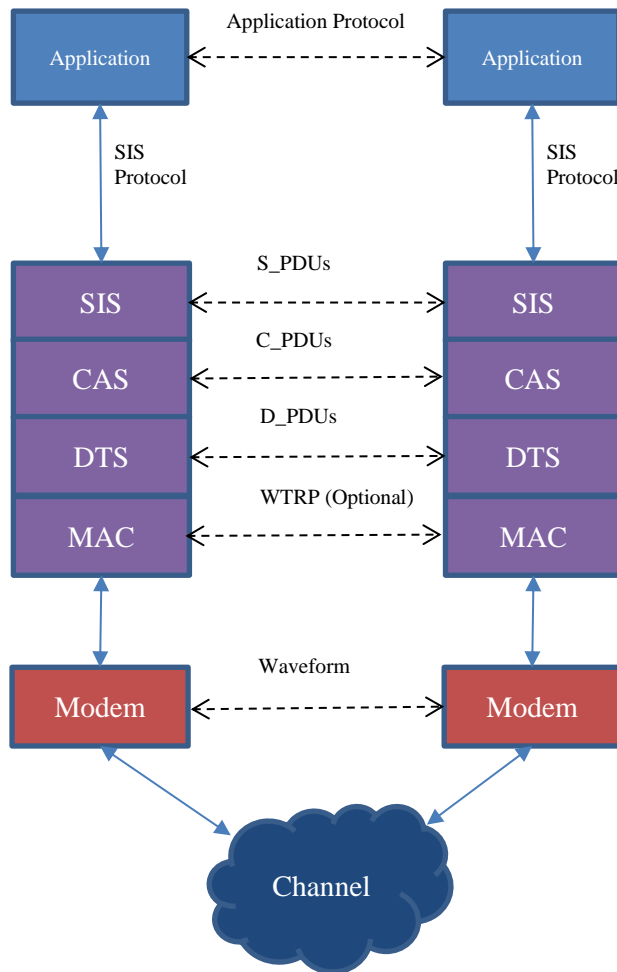### R.4.1. STANAG 5066 Ed3

Figure R-1 STANAG 5066 Ed3 Architecture

The STANAG 5066 Ed3 protocol architecture is illustrated above. Notes on this architecture.

1. Modems communicate with a waveform, such as STANAG 5069. There will be other hardware needed to connect a modem to an HF Channel (or to another channel such as UHF).
2. Communication between STANAG 5066 servers (shown in yellow) and modems is proprietary.
3. STANAG 5066 Applications communicate with STANAG 5066 servers using the SIS protocol.
4. Applications running over STANAG 5066 define their own end to end application protocols, such as COSS used for ACP 127 messaging.

5. STANAG 5066 is defined as four layers, with PDUs defined at each layer that (except for MAC) are transferred by the protocol layer below. These layers are:
    a. Subnet Interface Sublayer (SIS) which communicates S_PDUs.
    b. Channel Access Sublayer (CAS) which communicates C_PDUs.
    c. Data Transfer Sublayer (DTS) which communicates D_PDUs.
    d. Media Access Control (MAC) sublayer.   The communication depends on the MAC layer choice:
        i. The MAC layer is optional and is not required on a channel with just two nodes.
        ii. Annex K (CSMA).   The MAC layer is procedure only, and no protocol is exchanged.
        iii. WTRP. This exchanges protocol, but uses special D_PDUs (rather than having MAC layer PDUs).
        iv. TDMA.   There is a placeholder (Annex M) for this, but there is no specification.

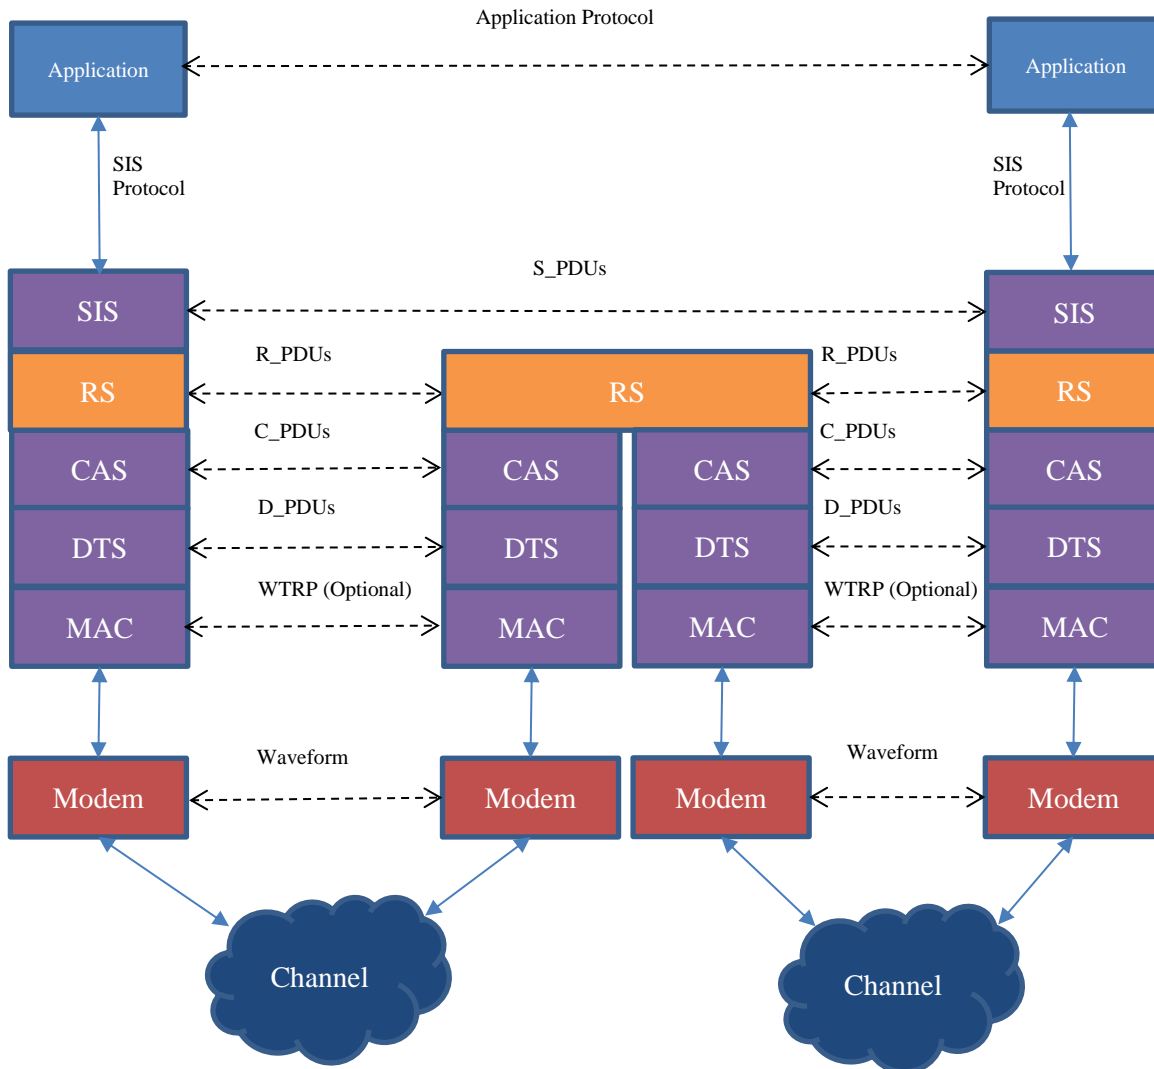## R.4.2. Routing with Multiple Independent Subnets

Figure R-2 Routing with Multiple Independent Subnets

Routing Sublayer (RS) is added as a STANAG 5066 layer between SIS and CAS, with exchange of R_PDUs. The SIS protocol is end to end, whereas CAS and the layers below are single hop.
The above diagram illustrates how the routing sublayer is used with two independent channels. If the LHS node in the above diagram sends data to the RHS node which is not connected, the routing sublayer communicates to the intermediate node the destination. This enables the middle node to route traffic between the two nodes. Note that this routing is completely transparent to the application.
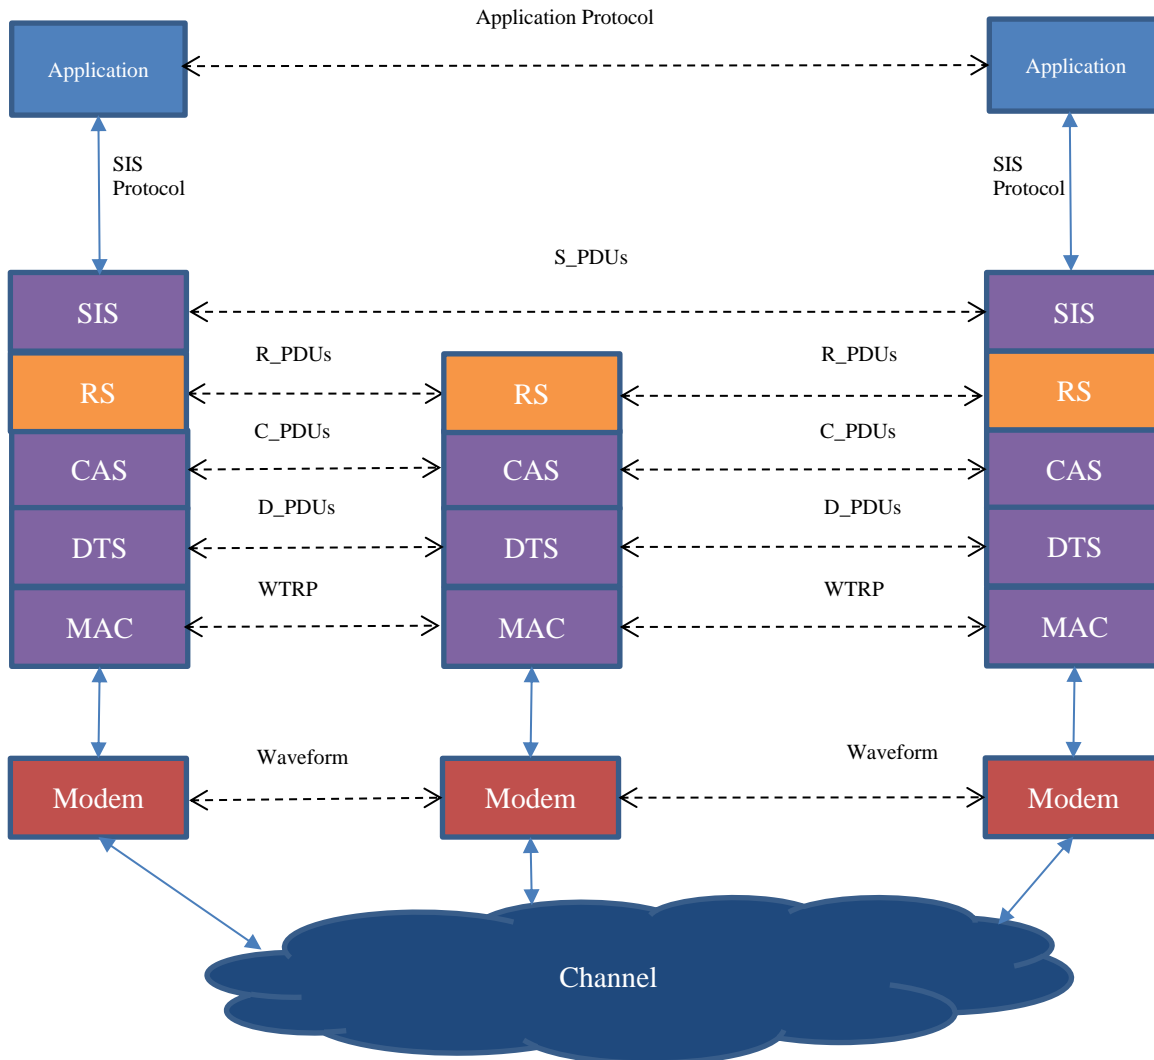
R.4.3. Routing with Single Channel

Figure R-3 – Routing with Single Channel

The above diagram shows how the routing sublayer works with a shared channel. WTRP is the only MAC layer that supports such a channel, so this description is given for WTRP, although the routing sublayer would work for a different MAC layer with the same characteristics.

Although all three modems are connected to the same channel, the diagram above relates to a scenario where the middle node can communicate with both of the end nodes, but traffic does not propagate between the end nodes. WTRP supports this configuration and will ensure that only one node transmits onto the channel at any one time.

All transmissions onto the channel are broadcast at the physical layer. When the LHS node transmits data to the RHS node, the address of the RHS node is included in the routing sublayer protocol, and the lower layers

send traffic to the middle node.   The middle node receives this message which is handled by the routing sublayer, which will send the data through the same stack but addressed to the RHS node. Both RHS and LHS nodes will hear data transmitted by the middle node, but will only handle traffic addressed to them.

R.4.4. Transparent Client Access to Multiple Subnets
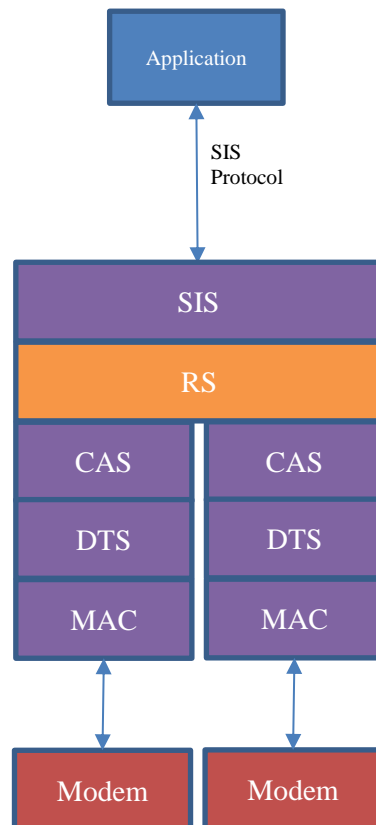

Figure R-4 – Client Access to Multiple Subnets

When a system is connected to two or more modems through independent STANAG 5066 stacks, the routing sublayer (RS) enables SIS provision to an application to transparently access the lower layers based on STANAG 5066 address of the destination.   The application does not need to be aware of which modem and stack is used.

R.5. Routing Layer PDUs

The routing sublayer uses five R_PDUs.   These use a single leading byte to identify the type of the PDU.

R.5.1. Direct R_PDU

	MSB							LSB
	7	6	5	4	3	2	1	0
0	Type = 0				Not used			
1	S_PDU							
n								

Figure R-5 – Direct R_PDU

The Direct R_PDU is used when the next hop is the final destination for the data. There is a single byte header of Type=0, and the rest of the data is the encapsulated S_PDU.

R.5.2. Indirect R_PDU

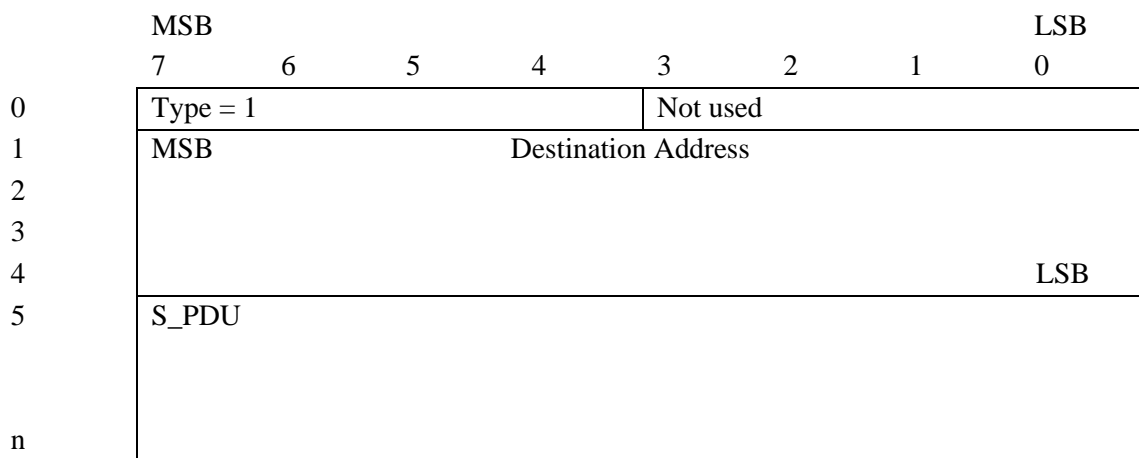|	MSB							LSB
	7	6	5	4	3	2	1	0
0	Type = 1				Not used			
1	MSB			Destination Address				
2								
3								
4								LSB
5	S_PDU							
n								

Figure R-6 Indirect R_PDU

The Indirect R_PDU is used when the next hop is not the final destination. The Indirect R_PDU is type 1, and includes the STANAG 5066 address of the destination. This will enable CAS and layers below to send the S_PDU to the next hop.

R.5.3. Broadcast R_PDU

|	MSB							LSB
	7	6	5	4	3	2	1	0
0	Type = 2				Not used			
1	S_PDU							
n								

Figure R-7 – Broadcast R_PDU

The Direct R_PDU is used when data is being broadcast (or multicast), for the initial broadcast when there is a single subnet There is a single byte header of Type=2, and the rest of the data is the encapsulated S_PDU.

R.5.4. Multicast R_PDU

	MSB							LSB
	7	6	5	4	3	2	1	0
0	Type = 3				Not used			
1	MSB			Original Sender				
2								
3								
4								LSB
5	MSB			Destination Count=d				LSB
6	MSB							
				Destination Address				
								LSB
9								
10	MSB							
				Destination Address				
								LSB
13								
5 + 4*d	S_PDU							
n								

Figure R-8 – Multicast R_PDU

The Multicast R_PDU is used when a broadcast/multicast is being relayed or for initial broadcast/multicast when there are multiple subnets. The STANAG 5066 address of the node that originally sent the broadcast is included, as this will not be available at on relay. Then there is a list of nodes to which the message needs to be sent. This is to ensure that a message is only re-broadcast when needed.

The number of destination STANAG 5066 addresses is encoded as Destination Count. This allows the list of destination addresses to be determined.

R.5.5. Routing Update R_PDU

| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |

MSB ... LSB

|								
0	Type = 4				Not used			
1	Routing update information							
n								

Figure R-9 – Routing Update R_PDU

The Routing Update R_PDU is used to share routing information between nodes on different subnetworks. This is type 2.   Routing update is a sequence of five byte information blocks encoded as follows.

MSB ... LSB

	7	6	5	4	3	2	1	0
0	MSB			Destination Address				
1								
2								
3								LSB
4	Hop Count							

Figure R-10 – Block Encoding of Routing Update R_PDU

The first four bytes of each routing update information block is the STANAG 5066 address of the node being reported on. The fifth byte is a Hop Count, which indicates the number of routing hops to this node from the node that is sending the routing update. Where the node is directly reachable from the reporting node, hop count is set to 1.

R.6. Procedures of Operation

This section describes how the routing sublayer works.

R.6.1. CAS Layer Interface

The interface to CAS layer is modified as follows:

1.   Where STANAG 5066 Annex B uses S_PDU, R_PDU is used when the routing sublayer is used.
2.   Other parameters to CAS layer are specified in the following procedures

R.6.2. Determining Routing

The routing sublayer needs to know the full set of addressable STANAG 5066 nodes and for each node it has to

determine the next hop.   This information can be obtained by four mechanisms:

1. Configuration.   For example when a subnet is set up with a specified set of STANAG 5066 nodes. This could be a fixed configuration or a configuration updated dynamically by an external process.
2. From Wireless Token Ring Protocol, as specified in Annex L, which supports topologies where nodes are not directly connected. Information on this topology is passed upwards using an Annex J M_ service primitive.   This is specified in Annex L as the next hop table which gives a list of all nodes in the ring, and indicates which nodes are directly connected.   For nodes that are not directly connected it indicates a preferred data relay node, which is the directly connected node that is recommended to use for data relay.   Annex L makes the next hop table available to the routing sublayer.
3. By discovery, using the mechanism specified by "HF Discovery, Ping and Traffic Load" (S5066-APP2).
4. Using the Routing Update procedures of this specification, which enables discovery of nodes which can be routed to indirectly over subnets that are connected indirectly.

R.6.3. Submission Procedure

The SIS layer will provide an S_PDU and other parameters to be passed to the lower layers.

1. If the target address is a broadcast address, there are three distinct scenarios:
   a. If the S_PDU is provided by the SIS layer and there is a single subnet, encapsulate the S_PDU in a Broadcast R_PDU. Then pass the R_PDU and submission parameters to CAS layer.
   b. If the S_PDU is provided by the SIS layer and there are multiple subnets for each of the connected subnets determine the list of destination addresses for each subnet, so that all target addresses are reached by exactly on subnet.   Then encapsulate the S_PDU in a Multicast S_PDU, with the Original Sender set to the local node address and the destination addresses included.
   c. If the S_PDU is provided by relay, there will be an Original Sender and a list of destination addresses to handle. Determine which subnet will handle each address, and for each subnet submit a Multicast R_PDU.

2. If the target address is a single node which does not have routing information, the SIS request is rejected with the appropriate C_UNIDATA_REJECT or C_EXPEDITED_UNIDATA_REJECT, with reason "Destination Node Not Responding".   This is the best reason code available. On review of STANAG 5066, it is recommended to consider adding in a new and more accurate reason code for this situation.

3. If the target address is a single node which can be directly reached on a connected subnet, then encapsulate the S_PDU in a Direct R_PDU and pass to CAS layer with the provided parameters.

4. If the target address is a single node which can be reached indirectly, then encapsulate the S_PDU in a Indirect R_PDU, with the destination address set to the target address. Then pass to the CAS layer the R_PDU with SIS submission parameters with the target address replaced by the next hop address.

R.6.4. Reception Procedure

The CAS layer will provide received R_PDUs to the routing sublayer.Handling of Routing Update R_PDUs is described in the next section.   The different PDUs are handled as follows:

1. The S_PDU from Direct R_PDUs are passed directly to the SIS layer.

2. When an Indirect R_PDU is received, the S_PDU is extracted. Then Submission Procedure is followed using the S_PDU, service parameters associated with the received R_PDU and the target address being the destination address from the Indirect R_PDU. This process is termed relaying.

3. When a Broadcast R_PDU is received over WTRP, the Relay Responsible List of nodes provided by the MAC layer is considered. For each of these addresses, determine the subnet to be used for relay. Use multicast submission procedure for this list of addresses, with the broadcast sender used as original sender.

4. When a Broadcast R_PDU is received over a subnet that is not WTRP, as list of nodes is generated from the broadcast/multicast address, with the local node and original sender removed, For each of these addresses, determine the subnet to be used for relay. Use multicast submission procedure for this list of addresses, with the broadcast sender used as original sender.

5. When a Multicast R_PDU is received, the following procedure is followed.

   a. If the original sender is the local node, the PDU is discarded an processing stops.
   b. The S_PDU is delivered locally following standard procedures.
   c. Consider the list of destination addresses included in the Multicast R_PDU. If the subnet is WTRP, consider the Relay Responsible List, and eliminate all addresses that are not in this list. If there are no addresses remaining, processing stops.
   d. Pass the PDU for relay using the multicast submission procedure with this list of addresses and the original sender.

In the event of any failure of a directly submitted R_PDU, errors are passed up to the SIS layer.
In the event of any failure of a relayed R_PDU the error must be handled by the routing sublayer. In the event of TTL expiry, the R_PDU is discarded. In the event of any other error, the R_PDU is resubmitted. The routing calculation must be repeated for the destination, as the preferred next hop may have changed.

R.6.5. Updating Routing

An node connected to multiple subnets must use the Routing Update R_PDU to communicate routing information to each subnet on the nodes that cannot be reached on that subnet, but can be reached on other subnets or can be reached with less hops on other subnets. This information is shared at intervals or whenever routing information changes.
On reception of a Routing Update R_PDU, the node must update its local configuration of reachable nodes.
The default approach to sharing Routing Update R_PDUs on a subnet is to use ARQ communication to each node on the subnet. This will ensure that full routing information is shared about all reachable nodes to all nodes.
For an Annex K (CSMA) network or an WTRP network where all nodes are directly reachable, Routing Update R_PDUs may be shared using non-ARQ broadcast.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

Annex S: SIS Access Protocol (mandatory)
DRAFT 1.0 of 9th April 2020

This annex defines a protocol for clients to access the STANAG 5066 Subnet Interface Service. This enables multiple independent clients to easily share and multiplex over a single channel.

S.1. Accessing the Subnet Interface Service

Figure S-1: SIS Access Protocol Model

A client-server relationship can be used to govern the interaction between the HF subnetwork and the users of the subnetwork, as shown in Figure S-1. The users (clients) request the services provided by the HF subnetwork (server). The service provided by the server is application independent and common to all clients irrespective of the task they may perform.

Clients are attached to the Subnetwork Interface Sublayer at Subnetwork Access Points (SAPs). There can be multiple clients simultaneously attached to the Subnetwork Interface Sublayer. Each SAP is identified by its SAP Identifier (SAP ID)[1]. The SAP ID is a number in the range 0-15; hence there can be a maximum of 16 clients attached to the Subnetwork Interface Sublayer of a single node.

The SIS Access protocol specified in this annex allows a client to access the Subnet Interface Service. There can be one client for each SAP, which enables multiplexing of multiple independent clients.

The SIS service specified in Annex A allows clients to access the SIS service by any mechanism. The mechanism specified in this annex is mandatory, to enable independent interoperable integration between STANAG 5066 severs and applications operating over STANAG 5066.
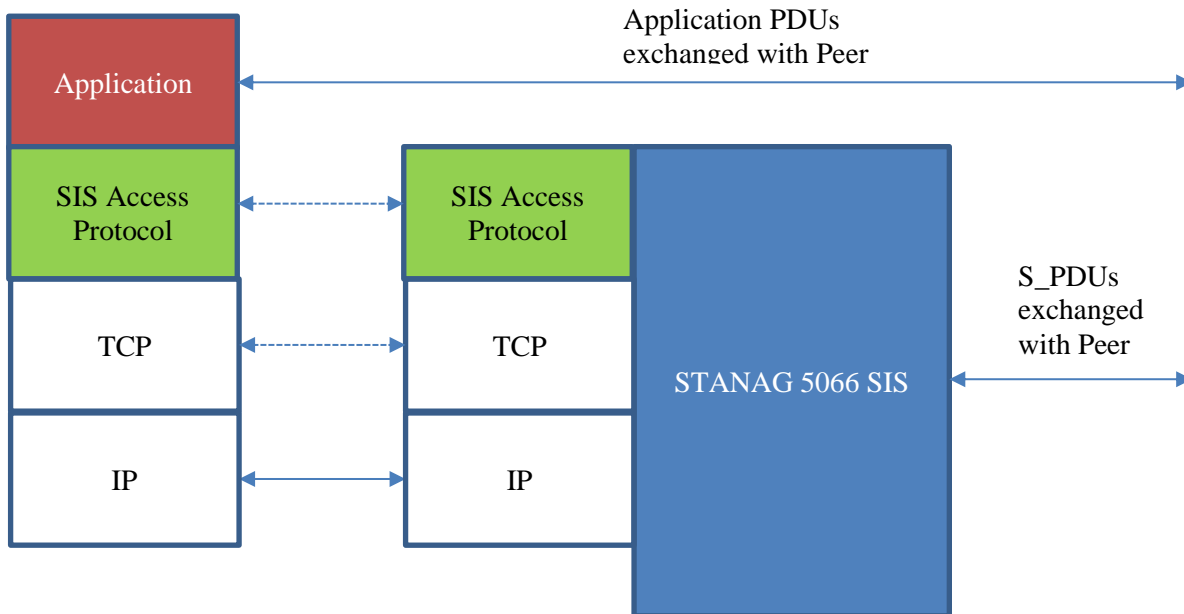
S.2. Mapping onto TCP

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

Figure S-2: SIS Access Protocol Mapping onto TCP

Figure S-2 shows how the SIS Access Protocol maps onto TCP/IP to connect an application to a SIS service. The SIS service communicates with its peer using S_PDUs. The application operating over SIS communicates with its peer using PDUs specific to the application. The SIS Access Protocol is a local protocol which does not have end to end interactions. It simply enables an application to interact with the local SIS service.

STANAG 5066 Annex A defines the core SIS services, which are extended by additional management and flow control services specified in this annex. Each of these service primitives is defined in a manner that is simple flow of data in one direction between application and SIS. This annex defines a framed encoding of each primitive, such that it can be mapped directly onto a TCP stream. This enables the application to access the SIS service over a TCP stream.

The Default SIS Port Number for the Raw SIS Socket Server **shall** be the decimal number '5066', a number registered with the Internet Assigned Number Authority (IANA) for this purpose. STANAG 5066 implementations, both clients and subnetworks, **should** be configurable to use any other valid port number as the default.

A SIS Access Protocol Server shall listen on the Default SIS Port Number for connection requests from clients. The Server shall accept sixteen connections simultaneously, so that every SAP can be served.

S_PRIMITIVES **shall** be sent over the SIS Access Protocol without any separating characters or framing data other than the generic S_PRIMITIVE encoding elements defined in this annex. No other messages between client and subnetwork **shall** be sent. Implementation-dependent communication between a client and subnetwork over the Raw SIS-Socket-Server Interface for the purposes of system management **may** be encapsulated within the S_MANAGEMENT_MSG_REQUEST and S_MANAGEMENT _MSG_INDICATION primitives defined in this

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

annex. [NB: This does not preclude implementation-dependent communication over another socket on a different port number, but such use is outside of the scope of this STANAG and is not recommended.]

S.3. Management & Flow Control Services

Table A-1. Management and Flow Control Primitives

CLIENT -> SUBNETWORK INTERFACE	SUBNETWORK INTERFACE -> CLIENT
	S_SUBNET_AVAILABILITY (Subnet Status, Reason)
	S_DATA_FLOW_ON( )
	S_DATA_FLOW_OFF ( )
S_MANAGEMNT _MSG_REQUEST (MSG TYPE, MSG BODY)	S_MANAGEMENT_MSG_INDICATION (MSG TYPE, MSG BODY)
S_KEEP_ALIVE ( )	S_KEEP_ALIVE ( )

Table A-1 summarizes primitives to specify management and flow control operations, which extend the core SIS service primitives specified in Annex A. The details of each of these primitives is set out below.

S.3.1. Interface Flow Control Primitives: S_DATA_FLOW_ON and S_DATA_FLOW_OFF

Name :

S_DATA_FLOW_ON
S_DATA_FLOW_OFF

Arguments :
NONE
Direction :
Subnetwork Interface-> Client
Description :
The S_DATA_FLOW_ON and S_DATA_FLOW_OFF primitives **shall** [1] be issued by the Subnetwork Interface Sublayer to control the transfer of U_PDUs submitted by a client.

On receipt of an _DATA_FLOW_OFF primitive, the client **shall** [2] cease transferring U_PDUs over the interface.

Transfer over the interface of U_PDUs by the client **shall** [3] be enabled following receipt of an S_DATA_FLOW_ON primitive.

The Subnetwork Interface Sublayer can use these two primitives (or other mechanisms) to control the flow of data from locally attached clients. U_PDUs from an attached client to which the S_DATA_FLOW_OFF primitive has been sent may be discarded by the Subnetwork Interface Sublayer without acknowledgement, indication, or warning.

A client **shall** [4] not control the flow of data *from* the subnetwork by any mechanism, explicit or

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

implicit.

All clients **shall** [5] be ready to accept at all times data received by the HF Node to which it is bound; clients not following this rule may be disconnected by the node.

S.3.2.   S_MANAGEMENT_MSG_REQUEST Primitive

Name :
S_MANAGEMENT_MSG_REQUEST
Arguments :
1. MSG TYPE
2. MSG BODY
Direction :
Client-> Subnet Interface
Description :
The S_MANAGEMENT_MSG_REQUEST primitive **shall** [1] be issued by a client to submit a "Management" message to the Subnetwork.

The complex argument MSG may be implementation dependent and is not specified in this version of STANAG 5066. At present, a minimally compliant HF subnetwork implementation **shall** [2] be capable of receiving this primitive, without further requirement to process its contents.

Depending on the value of the complex argument *MSG*, this primitive can take the form of a Command (e.g. Go-To-EMCON, Go-Off-Air, etc.) or of a Request (e.g. Request-For-Subnetwork-Statistics, Request-For-Connected-client-Information, etc.).

Note that this primitive is not intended to allow for the transmission of management coordination messages over the air. This is an interaction between peer subnet management clients and as such shall be accomplished using the UNIDATA primitive defined elsewhere in this annex.

S.3.3.   S_MANAGEMENT_MSG_INDICATION Primitive

Name :
S_MANAGEMENT_MSG_INDICATION
Arguments :
1. MSG TYPE
2. MSG BODY
Direction :
Subnetwork Interface-> Client
Description :
The S_MANAGEMENT_MSG_INDICATION primitive **shall** [1] be issued by the Subnetwork to send a "Management" message to a client.

The complex argument MSG may be implementation dependent and is not specified in this version of STANAG 5066. At present, a minimally compliant client **shall** [2] be capable of receiving this primitive, without further requirement to process its contents.

As implementation options, the complex argument *MSG* could take several values such as: Subnetwork-Statistics, Connected-client-Information, etc. This primitive could be issued either in response to a

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S_MANAGEMENT_MSG_REQUEST or asynchronously by the Subnetwork.

S.3.4.   S_KEEP_ALIVE Primitive

Name :
S_KEEP_ALIVE
Arguments :
NONE
Direction :
Client-> Subnetwork Interface
Subnetwork Interface-> Client
Description :
The S_KEEP_ALIVE primitive can be issued as required (e.g. during periods of inactivity) by the clients and/or the Subnetwork Interface to sense whether the physical connection between the client and the Subnetwork is alive or broken. This primitive may be redundant if the implementation of the physical connection provides an implicit mechanism for sensing the status of the connection.

A minimally compliant implementation of a client or subnetwork interface is not required to generate the S_KEEP_ALIVE primitive except in response to the receipt of an S_KEEP_ALIVE primitive.

When the S_KEEP_ALIVE Primitive is received, the recipient (i.e, client or Subnetwork Interface) **shall** respond with the same primitive within 10 seconds.

If a reply is not sent within 10 seconds, no reply **shall** [2] be sent.

A client or Subnetwork Interface **shall** [3] not send the S_KEEP_ALIVE Primitive more frequently than once every 120 seconds to the same destination.

S.3.5.   S_SUBNET_AVAILABILITY Primitive

Name :
S_SUBNET_AVAILABILITY
Arguments :
1. Node Status
2. Reason
Direction :
Subnetwork Interface-> Client

Description:
The S_SUBNET_AVAILABILITY primitive may be sent asynchronously to all or selected clients connected to the Subnetwork Interface Sublayer to inform them of changes in the status of the node to which they are attached. For example, clients can be informed using this primitive that available resources (e.g., bandwidth) have been temporarily reserved for other clients. Alternatively, this primitive could be used to inform clients that the node has entered an EMCON state and as a result they should only expect to receive Data and will not be allowed to transmit data.

The contents of this primitive are implementation dependent and not specified in this version of STANAG 5066. At present, a minimally compliant client implementation **shall** [1] be capable of receiving this primitive, without further requirement to process its contents.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

As implementation options, the *Node Status* argument could specify the new Status of the node. Possible values of this argument could be ON, OFF, Receive-Only, Transmit-Only-to-Specific-Destination-Node/SAP, etc.

Node Status	Value
OFF	0
ON	1
Receive-Only	2
Half-Duplex	3
Full-Duplex	4
Transmit-Only	5

If the Subnetwork Status is other than ON, the *Reason* argument explains why. Values of this argument shall be as specified below.

The value assigned to each Node Status **shall** be used to represent the reason in SIS Access Protocol (Annex S).

Reason	Value
unspecified	0
Local Node in EMCON	1
Reserved	2

The value assigned to each reason **shall** be used to represent the reason in SIS Access Protocol (Annex S).

S.4. Encoding of Primitives

The encoding of the S_Primitives for communication using the protocol specified in this annex **shall** [(1)] be in accordance with text and figures in the subsections below.

S.4.1. Generic Field Encoding Requirements

Unless noted otherwise, the bit representation for argument values in an S_Primitive **shall** [(1)] be encoded into their corresponding fields in accordance with CCITT V.42, 8.1.2.3, which states that:
- when a field is contained within a single octet (i.e, eight bit group), the lowest bit number of the field **shall** [(2)] represent the lowest-order (i.e., least-significant-bit) value;
- | when a field spans more than one octet, the order of bit values within each octet **shall** [(3)] progressively decrease as the octet number increases. The lowest bit number associated with the field represents the lowest-order value.

The 4-byte address field in the S_primitives **shall** [(4)] carry the 3.5-byte address and address-size information defined in A.2.2.28.1. The lowest order bit of the address shall be placed in the lowest order bit position of the field (generally bit 0 of the highest byte number of the field), consistent with the mapping specified in Section C.3.2.6 for D_PDUs.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.4.2.   S_Primitive Generic Elements and Format

As shown in Figure A-1(a), all primitives **shall** [(1)] be encoded as the following sequence of elements:

- a two-byte S_Primitive preamble field, whose value is specified by the 16-bit Maury-Styles sequence below;

- a one-byte version-number field;

- a two-byte Size_of_Primitive field;

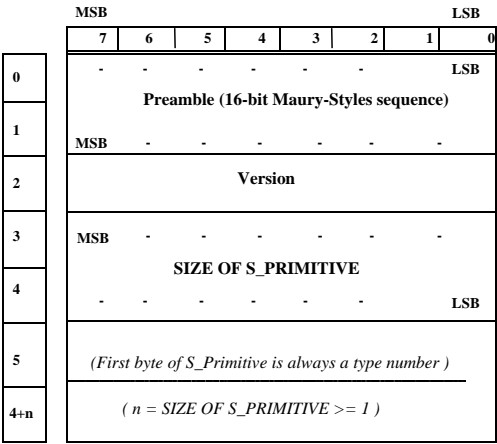- a multi-byte field that contains the encoded S_Primitive.

	MSB							LSB
	7	6	5	4	3	2	1	0
0	-	-	-	-	-	-		LSB
	Preamble (16-bit Maury-Styles sequence)							
1	MSB	-	-	-	-	-	-	
2	Version							
3	MSB	-	-	-	-	-	-	
	SIZE OF S_PRIMITIVE							
4	-	-	-	-	-	-		LSB
5	*(First byte of S_Primitive is always a type number)*							
4+n	*(n = SIZE OF S_PRIMITIVE >= 1)*							

Figure S-3: Element-Sequence *Encoding of "S_" Primitives*

The S_Primitive preamble field **shall** [(2)] be encoded as the 16-bit Maury-Styles sequence shown below, with the least significant bit (LSB) transmitted first over the interface:

(MSB) 1 1 1 0 1 0 1 1 1 0 0 1 0 0 0 0 (LSB)

i.e., with the multi-byte S_Primitive field represented in hexadecimal form as 0xEB90, the least-significant bits of the sequence **shall** [(3)] be encoded in the first byte (i.e, byte number 0) of the preamble field and the most significant bits of the sequence **shall** [(4)] be encoded in the second byte (i.e, byte number 1) of the preamble field as follows:

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

	MSB 7	6	5	4	3	2	1	LSB 0	
0	1	0	0	1	0	0	0	0	0x90
1	1	1	1	0	1	0	1	1	0xEB

Figure S-4: Encoding of Maury-Styles Preamble-Sequence in "S_" Primitives

[*Note*: This encoding of the Maury-Styles preamble sequence is an exception to the general requirement of section S.4.1 for field encoding.]

Following the Maury-Styles sequence, the next 8 bit (1-byte) field **shall** [5] encode the 5066 version number. For this version of STANAG 5066, the version number **shall** [6] be all zeros, i.e, the hexadecimal value 0x00, as shown in Figure S-3.

The next 16 bit (two-byte) field **shall** [7] encode the size in bytes of the S_primitive-dependent field to follow, exclusive of the Maury-Styles sequence, version field, and this size field. The LSB of the of the size value **shall** be mapped into the low order bit of the low-order byte of the field as shown in Figure S-3.

Unless specified otherwise, the order of bit transmission for each byte in the encoded S_Primitive **shall** [9] be as described in CCITT V.42 paragraph 8.1.2.2, which specifies the least significant bit (LSB, bit 0 in the figures below) of byte 0 **shall** [10] be transmitted first.

The sixth byte (i.e., byte number 5) of the sequence **shall** [11] be the first byte of the encoded primitive and **shall** be equal to the S_Primitive type number, with values encoded in accordance with the respective section that follows for each S_primitive

The remaining bytes, if any, in the S_Primitive **shall** [13] be transmitted sequentially, also beginning with the LSB of each byte, in accordance with the respective section that follows for each S_primitive.

In the subsections that follow, any bits in a S_Primitive that are specified as NOT USED **shall** [13] be encoded with the value "0" unless specified otherwise for the specific S_Primitive being defined.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.4.3.   S_BIND_REQUEST Encoding

The S_BIND_REQUEST primitive **shall** [1] be encoded as a four-byte field as follows:


Figure S-5: Encoding of S_BIND_REQUEST Primitive

The S_BIND_REQUEST SERVICE-TYPE field **shall** [2] be encoded as five subfields as follows:
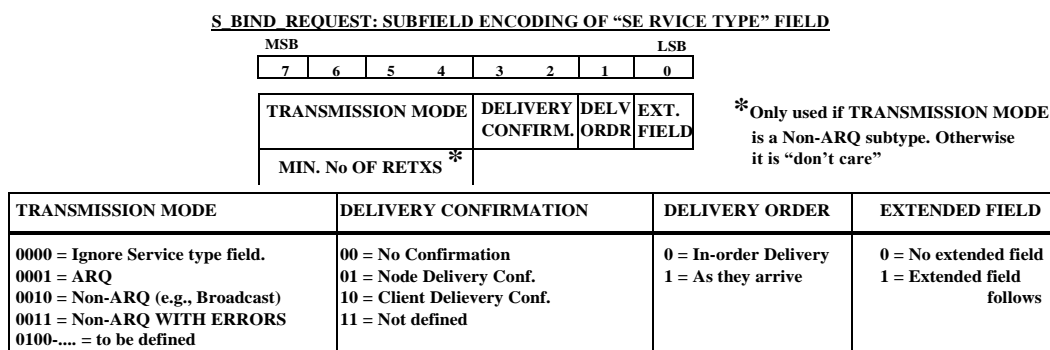

Figure S-6: Sub-field Attribute Encoding of S_BIND_REQUEST SERVICE-TYPE field.

Argument : SERVICE TYPE
Primitive : S_BIND_REQUEST

The SERVICE TYPE argument **shall** [3] specify the default type of service requested by the client. This type of service **shall** [4] apply to any U_PDU submitted by the client until the client unbinds itself from the node, unless overridden by the DELIVERY MODE argument of the U_PDU. A client **shall** [5] change the default service type only by unbinding and binding again with a new S_BIND_REQUEST.

The RANK value was present in Edition 3 of STANAG 5066, but is not used in this edition. It should be set to zero on transmission and ignored on reception.

The SERVICE TYPE argument is complex, consisting of a number of attributes encoded as sub-fields. Although the exact number of attributes and their encoding is left for future definition and enhancement using the Extended Field attribute, the following attributes are mandatory:

1. *Transmission Mode for the Service.* --- ARQ or Non-ARQ Transmission Mode **shall**[6] be specified, with one of the Non-ARQ submodes if Non-ARQ was requested. A value of "0" for this attribute **shall**[7] be invalid for the SERVICE TYPE argument when binding. Non-ARQ transmission can have submodes such as: *Error-Free-Only* delivery to destination client, delivery to destination client even

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

with *some* errors.

2. *Data Delivery Confirmation for the Service* --- The client **shall** [8] request one of the Data Delivery Confirmation modes for the service. There are three types of data delivery confirmation:
     - None
     - Node-to-Node Delivery Confirmation
     - Client-to-Client Delivery Confirmation

The client can request explicit confirmation, i.e, Node-to-Node or Client-to-Client, from the Subnetwork to provide indication that its U_PDUs have been properly delivered to their destination. Explicit delivery confirmation **shall** [9] be requested only in combination with ARQ delivery.

[Note: The Node-to-Node Delivery Confirmation does not require any explicit peer-to-peer communication between the Subnetwork Interface Sublayers and hence it does not introduce extra overhead. It simply uses the

ACK (ARQ) confirmation provided by the Data Transfer Sublayer. Client-to-Client Delivery Confirmation requires explicit peer-to-peer communication between the Sublayers and therefore introduces overhead. It should be used only when it is absolutely critical for the client to know whether or not its data was delivered to the destination client (which may, for instance, be disconnected).]

3. *Order of delivery of any U_PDU to the receiving client.* --- A client **shall** [10] request that its U_PDUs are delivered to the destination client "in-order" (as they are submitted) or in the order they are received by the destination node.

4. *Extended Field* --- Denotes if additional fields in the SERVICE TYPE argument are following; at present this capability of the SERVICE TYPE is undefined, and the value of the Extended Field Attribute **shall** [11] be set to "0".

5. *Minimum Number of Retransmissions* --- This argument **shall** [12] be valid if and only if the Transmission Mode is a Non-ARQ type. If the Transmission Mode is a Non-ARQ type, then the subnetwork **shall** [13] retransmit each U_PDU the number of times specified by this argument. This argument may be "0", in which case the U_PDU is sent only once.
     [Note: In non-ARQ Mode, automatic retransmission a minimum number of times may be used to improve the reliability of broadcast transmissions where a return link from the receiver is unavailable for explicit retransmission requests.]

S.4.4. S_UNBIND_REQUEST Encoding

The S_UNBIND_REQUEST primitive **shall** [1] be encoded as a one-byte field as follows:

MSB							LSB
7	6	5	4	3	2	1	0

5 | TYPE *   TYPE = 2 (S_UNBIND_REQUEST)

Figure S-7: Encoding of S_UNBIND_REQUEST Primitive

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.4.5.   S_BIND_ACCEPTED Encoding

The S_BIND_ACCEPTED primitive **shall** [1] be encoded as a four-byte field as follows:

Figure S-8: Encoding of S_BIND_ACCEPTED Primitive

S.4.6.   S_BIND_REJECTED Encoding

The S_BIND_REJECTED primitive **shall** [1] be encoded as a two-byte field as follows:

Figure S-9: Encoding of S_BIND_REJECTED Primitive

S.4.7.   S_UNBIND_INDICATION Encoding

The S_UNBIND_INDICATION primitive **shall** [1] be encoded as a two-byte field as follows:

Figure S-10: Encoding of S_UNBIND_INDICATION Primitives

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.4.8.   S_SUBNET_AVAILABILITY Encoding

The S_SUBNET_AVAILABILITY primitive **shall** [(1)] be encoded as a three-byte field as follows:


Figure S-11: Encoding of S_SUBNET_AVAILABILITY Primitives.

The encoding of the NODE STATUS and REASON fields is implementation dependent.

S.4.9.   S_DATA_FLOW_ON and S_DATA_FLOW_OFF Encoding

The S_DATA_FLOW_ON and S_DATA_FLOW_OFF primitives **shall** [(1)] be encoded as one-byte fields as follows:

Figure S-12 Encoding of S_DATA_FLOW_ON and S_DATA_FLOW_OFF Primitives.

S.4.10.   S_KEEP_ALIVE Encoding

The S_KEEP_ALIVE primitive **shall** [(1)] be encoded as a one-byte field as follows:

Figure S-13: Encoding of S_DATA_FLOW_ON and S_DATA_FLOW_OFF Primitives.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.4.11.   S_MANAGEMENT_MSG_REQUEST and S_MANAGEMENT_ MSG_INDICATION Encoding

The S_MANAGEMENT_MSG_REQUEST and S_MANAGEMENT_MSG_ INDICATION primitives **shall** [1] be encoded as implementation-dependent variable-length fields as follows:


Figure S-14: Encoding of S_MANAGEMENT_MSG_REQUEST and S_MANAGEMENT_MSG_INDICATION Primitives.

The encoding of the MSG TYPE and MSG BODY fields is implementation dependent.

S.4.12.   S_UNIDATA_REQUEST Encoding

The S_UNIDATA_REQUEST primitive **shall** [1] be encoded as a variable-length field as follows:

Figure S-15: Encoding of S_UNIDATA_REQUEST Primitives.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

The SOURCE NODE ADDRESS and DESTINATION NODE ADDRESS fields **shall** [2] be encoded as specified in Section A.2.2.28.1.

The DELIVERY MODE field **shall** [3] be encoded as specified in Section S.4.16.2.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.4.13.   S_UNIDATA_INDICATION Encoding

The S_UNIDATA_INDICATION primitive **shall** [1] be encoded as a variable-length field as follows:
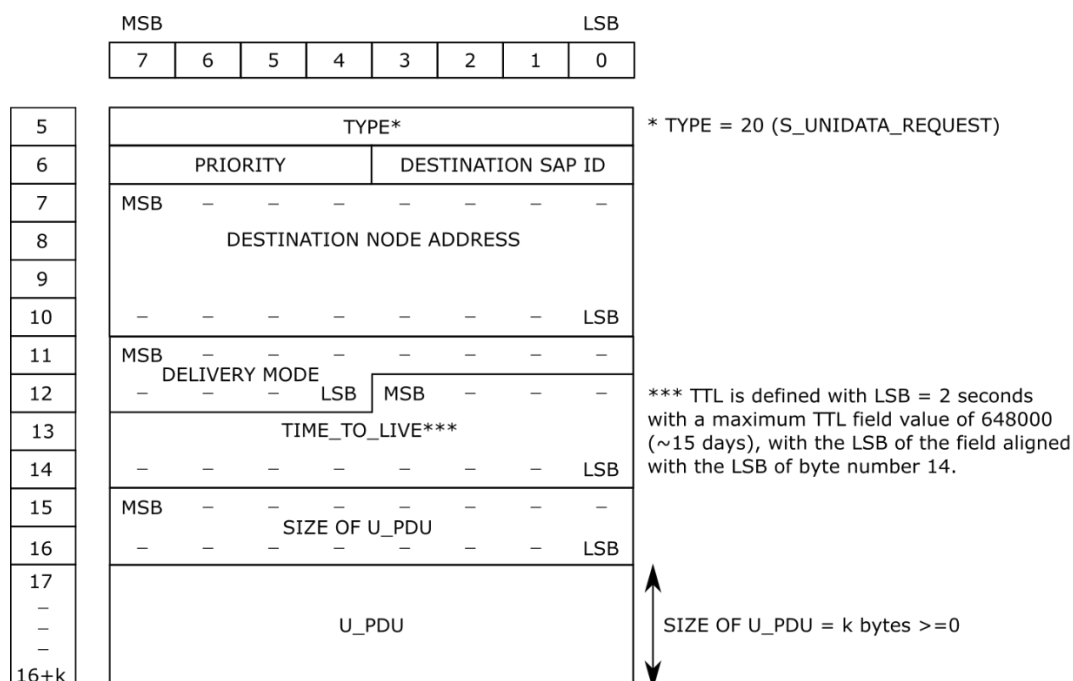

Figure S-16: Encoding of S_UNIDATA_INDICATION Primitives

The SOURCE NODE ADDRESS and DESTINATION NODE ADDRESS fields **shall** [2] be encoded as specified in Section S.4.16.1.

The TRANSMISSION MODE field **shall** [3] be encoded as specified in Section S.4.16.3.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.4.14.   S_UNIDATA_REQUEST_CONFIRM Encoding

The S_UNIDATA_REQUEST_CONFIRM primitive **shall** [(1)] be encoded as a variable-length field as follows:
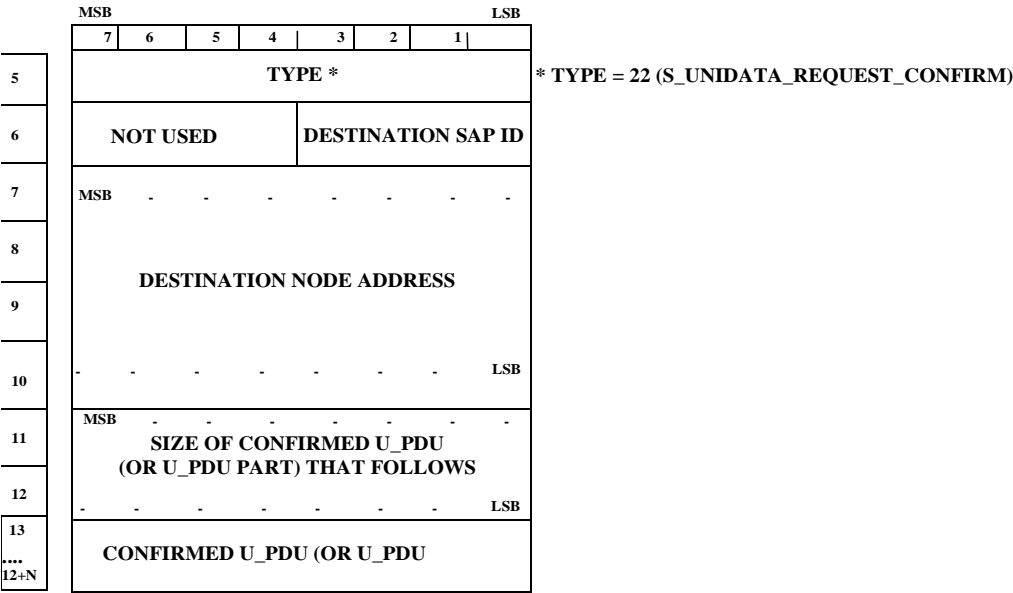

Figure S-17: Encoding of S_UNIDATA_REQUEST_CONFIRM Primitives.

The DESTINATION NODE ADDRESS field **shall** [(2)] be encoded as specified in Section S.4.16.1.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.4.15.   S_UNIDATA_REQUEST_REJECTED Encoding

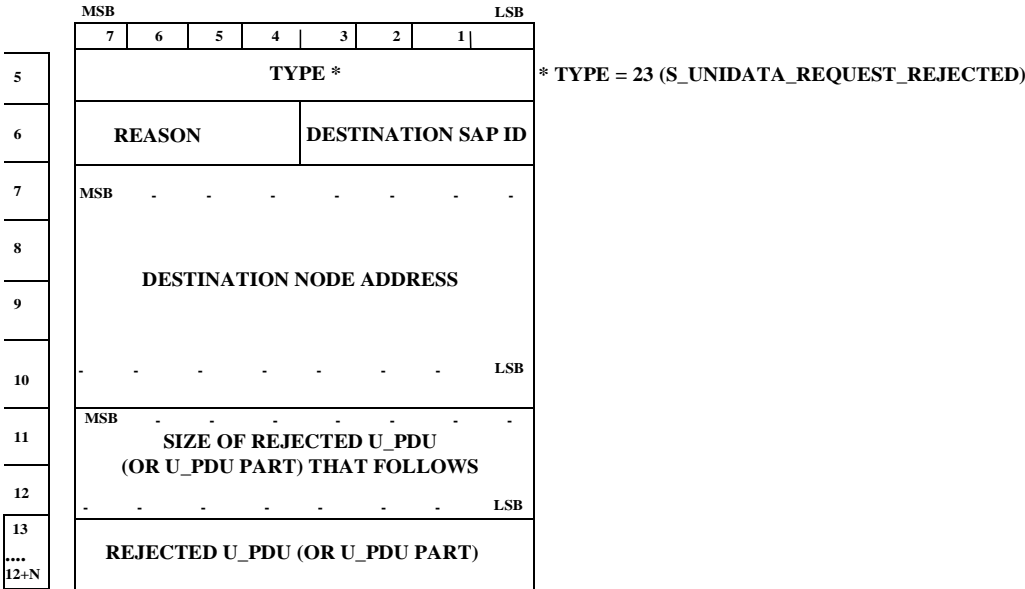The S_UNIDATA_REQUEST_REJECTED primitive **shall** [(1)] be encoded as a variable-length field as follows:

MSB							**LSB**	
	7	6	5	4	3	2	1	
5	TYPE *							* TYPE = 23 (S_UNIDATA_REQUEST_REJECTED)
6	REASON				DESTINATION SAP ID			
7	MSB							
8	DESTINATION NODE ADDRESS							
9								
10	. . . . LSB							
11	MSB . . . . SIZE OF REJECTED U_PDU (OR U_PDU PART) THAT FOLLOWS							
12	. . . . LSB							
13 .... 12+N	REJECTED U_PDU (OR U_PDU PART)							

Figure S-18: Encoding of S_UNIDATA_REQUEST_REJECTED Primitives.

The DESTINATION NODE ADDRESS field **shall** [(2)] be encoded as specified in Section S.4.16.1.

S.4.16.   Additional S_Primitive Encoding Requirements: Encoding of Common Fields

In order to clarify some of the procedures and tasks executed by the sublayers, additional details concerning some of the arguments of the Primitives described in previous sections are provided below.

S.4.16.1.   Node ADDRESS Encoding for all Primitives

Arguments :   SOURCE NODE ADDRESS, DESTINATION NODE ADDRESS, or REMOTE NODE ADDRESS
Primitives :   ALL "UNIDATA" primitives.

For reduced overhead in transmission, node addresses **shall** [(1)] be encoded in one of several formats that are multiples of 4-bits ("half-bytes") in length, as specified in Figure S-19.

Addresses that are encoded as Group node addresses **shall** [(2)] only be specified as the Destination Node address of Non-ARQ PDUs.

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

Destination SAP IDs and destination node addresses of ARQ PDUs and source SAP IDs and source node addresses of all PDUs **shall** [(3)] be individual SAP IDs and individual node addresses respectively.

ENCODING FORMAT OF ADDRESS FIELDS USED IN THE "S_" PRIMITIVES"
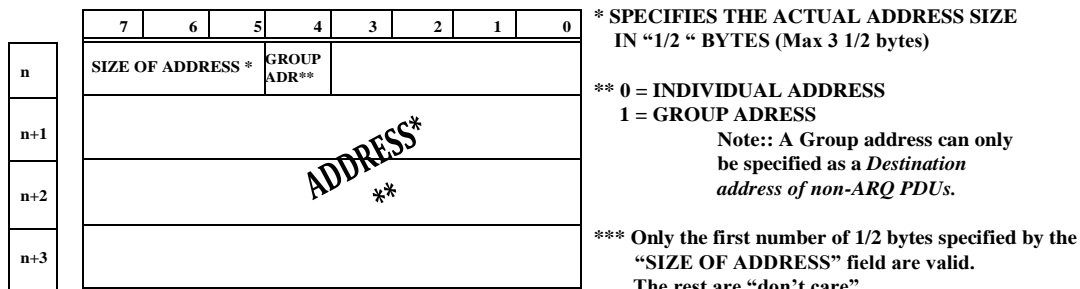

Figure S-19: Encoding of Address Fields in S_Primitives.

S.4.16.2.   Delivery-Mode Encoding for the S_UNIDATA_REQUEST Primitive

Argument : DELIVERY MODE
Primitive : S_UNIDATA_REQUEST , S_EXPEDITED_UNIDATA_REQUEST

The DELIVERY MODE is a complex argument consisting of a number of attributes, as specified here. The DELIVERY MODE argument **shall** [(1)] be encoded as shown in Figure S-20.

The value of the DELIVERY MODE argument can be "DEFAULT", as encoded by the Transmission Mode attribute. With a value of "DEFAULT", the delivery mode for this U_PDU **shall** [(2)] be the delivery mode specified in the *Service Type* argument of the S_BIND_REQUEST. A non-DEFAULT value **shall** [(3)] override the default settings of the Service Type for this U_ PDU.

The attributes of this argument are similar to those described in the *Service Type* argument of the S_BIND_REQUEST:

6. *Transmission Mode of this U_PDU.* --- ARQ or Non-ARQ Transmission can be requested. A value of "0" for this attribute **shall** [(4)] equal the value "DEFAULT" for the Delivery Mode. If the DELIVERY MODE is "DEFAULT", all other attributes encoded in the argument **shall** [(5)] be ignored.
7. *Data Delivery Confirmation for this PDU* --- None, node-to-node, or client-to-client.

8. *Order of delivery of this PDU to the receiving client.* --- A client may request that its U_PDUs are delivered to the destination client "in-order" (as they are submitted) or in the order they are received by the destination node.
9. *Extended Field* --- Denotes if additional fields in the DELIVERY MODE argument are following; at present this capability of the DELIVERY MODE is undefined, and the value of the Extended Field

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

Attribute **shall** [(6)] be set to "0".

10. *Minimum Number of Retransmissions* --- This argument **shall** [(7)] be valid if and only if the Transmission Mode is a Non-ARQ type or sub-type. If the Transmission Mode is a Non-ARQ type or subtype, then the subnetwork **shall** [(8)] retransmit each U_PDU the number of times specified by this argument. This argument may be "0", in which case the U_PDU is sent only once.

> [Note: In non-ARQ Mode, automatic retransmission a minimum number of times may be used to improve the reliability of broadcast transmissions where a return link from the receiver is unavailable for explicit retransmission requests.]

ENCODING OF "DELIVERY MODE" FIELD OF S_UNIDATA_REQUEST AND S_EXPEDITED_UNIDATA_REQUEST

MSB							LSB
7	6	5	4	3	2	1	0

TRANSMISSION MODE				DELIVERY CONFIRM.	DELV ORDR	EXT. FIELD

| MIN. No OF RETXS * | | | | | | |

***Only used if TRANSMISSION MODE is a Non-ARQ subtype. Otherwise it is "don't care"**

TRANSMISSION MODE	DELIVERY CONFIRMATION	DELIVERY ORDER	EXTENDED FIELD
0000 = Ignore Service type field. 0001 = ARQ 0010 = Non-ARQ (e.g., Broadcast) 0011 = Non-ARQ WITH ERRORS 0100-.... = to be defined	00 = No Confirmation 01 = Node Delivery Conf. 10 = Client Delievery Conf. 11 = Not defined	0 = In-order Delivery 1 = As they arrive	0 = No extended field 1 = Extended field follows

Figure S-20: Encoding of the Delivery Mode field in the S_UNIDATA_REQUEST and primitives

S.4.16.3. TRANSMISSION-MODE Encoding for the S_UNIDATA_INDICATION Primitive

Argument: TRANSMISSION-MODE
S_Primitives: S_UNIDATA_INDICATION, S_EXPEDITED_UNIDATA_INDICATION

ENCODING OF "TRANSMISSION MODE" FIELD:
S_UNIDATA_INDICATION AND S_EXPEDITED_UNIDATA_INDICATION

MSB							LSB
7	6	5	4	3	2	1	0

TRANSMISSION MODE							

TRANSMISSION MODE
0000 = Not Used 0001 = ARQ 0010 = Non-ARQ (Broadcast) 0011 = Non-ARQ WITH ERRORS 0100 ....= to be defined

Figure S-21: Encoding of Transmission Mode Field in S_UNIDATA_INDICATION primitive.

The subnetwork notifies a client of the transmission-mode used to deliver a U_PDU or Expedited U_PDU Argument with the TRANSMISSION-MODE argument. The TRANSMISSION-MODE argument in the S_UNIDATA_INDICATION and S_EXPEDITED_UNIDATA_INDICATION Primitives **shall** [(1)] be encoded as shown in Figure S-21.

[Note: The unused bits in this argument are allocated to the SOURCE SAP_ID argument encoding for both the S_UNIDATA_INDICATION and S_EXPEDITED_UNIDATA_INDICATION Primitives.]

NATO UNCLASSIFIED
Releasable to the Public
DRAFT FOR REVIEW

ANNEX S
STANAG 5066
Edition 4

S.5. Changes in Edition 4

This Annex is built from information in Edition 3 Annexes A and F. It supports the services defined in Annex A, and the protocol specified is unchanged from Edition 3.

Annex T: STANAG 5066 TRANSEC Crypto Sublayer using AES and other Protocols (Optional)

DRAFT 1.0 of 2nd April 2020

The TRANSEC Crypto Sublayer is an optional sublayer that provides TRANSEC. It defines a framework for using arbitrary stream cryptography to provide TRANSEC. It specifies how to use Advanced Encryption Standard (AES) in this framework.

The TRANSEC specified in Annex T can be used as an alternative to synchronous serial crypto devices connected to the DTS using STANAG 5066 Annex D.

This annex is new to Edition 4 of STANAG 5066.

T.1. Overview

This annex specifies protocol for supporting a Crypto Layer between STANAG 5066 and Modem, following the STANAG 5066 model. This protocol defines a generic framework for use with different encryption algorithms, with a specific mapping for AES encryption.


Figure T-1: STANAG 5066 TRANSEC Crypto Architecture

The STANAG 5066 architecture for use of TRANSEC Crypto is shown in Figure T-1. Crypto is used immediately above the modem, which is the lowest possible point that Crypto could be inserted. A stream crypto is used, so that modem data errors only impact that data. When crypto is synchronized, it does not introduce additional errors. This is an efficient approach.

HF traffic is easy to monitor and even the lowest layers of STANAG 5066 contain information that is of potential interest. Therefore, it makes good security sense to perform encryption at this lowest possible layer.

T.1.1. Benefits of a Protocol Approach

This specification introduces a protocol approach, which gives the benefits of the architecture and removes the

overheads and issues associated with use of sync serial following Annex D.   A number of implementation approaches are possible.   An implementation approach that could be taken is:

1. Use the open TCP protocol interface specified in MIL-STD-188-110D Annex A to communicate between Crypto Layer and modem.
2. Implement the Crypto Layer framing as specified here, as part of the STANAG 5066 server.   This can offer:
    a. Built in AES support.
    b. Plugin options to drive other Crypto.

T.2. Crypto Considerations

This section considers a number of crypto issues.

T.2.1. AES

AES (Advanced Encryption Standard) is a widely adopted US Government standard for encryption.   It is widely used for commercial, government and military operation. .Therefore, use of AES is specified in this annex.

T.2.2. COMSEC and TRANSEC

TRANSEC (Transmission Security) is the protection applied at the lowest communication level, of which Crypto between modem and STANAG 5066 is a good example.   Technically, this annex specifies TRANSEC.

COMSEC (Communication Security) is the protection applied to user data. TRANSEC **may** be used to provide COMSEC.   The model for use of Annex D is that COMSEC is provided by TRANSEC.   This annex defines TRANSEC.   It is a deployment decision as to whether this is also used for COMSEC.

T.3. Modem Service Specification

In order to understand how the Crypto Layer works, it is important consider the modem service interface.

T.3.1. Modem Transmission

Figure T-2: Transmission of D_PDUs to a Modem

STANAG 5066 will transmit over the modem a "bounded stream" of D_PDUs as shown in Figure T-2.   This is sent as a sequence of bytes, with request to start transmission implicit in the first byte.   When the last byte of

the last D_PDU is sent, it is marked as "end of stream".   The modem may pad after the last D_PDU in order to fill an exact number of blocks.

STANAG 5066 will also set parameters for the transmission, which will be fixed for the period of transmission and will deal with any errors reported. Parameters include: speed; interleaver; waveform; waveform-specific parameters; bandwidth.

For duplex and broadcast, transmissions may be of arbitrary length.   In other cases, transmissions are limited to 127.5 seconds and each D_PDU is marked with remaining transmission time (EOT) in units of 0.5 seconds.

This information is transparent to the modem, but important for overall operation. STANAG 5066 will know the speed of modem in order to calculate data to send and to keep up with the modem.   Transmission length is determined as start of transmission. However, it is desirable to defer choice of which data to send, to enable insertion of high priority data arriving after a transmission starts.

T.3.2. Modem Reception

Under good conditions, modem reception of data will be entirely symmetrical to transmission, and a bounded stream of D_PDUs will be provide to the receiving server.   On reception with modern waveforms, most transmission parameters are determined from the transmission.   Dealing with modem reception issues is key to Crypto Layer design.

T.3.3. Determining Transmission End

Data is often lost or corrupted during transmission.   A key problem is to determine end of transmission.   There are strict rules for modem transmission.   A receiving modem needs to apply heuristics to determine transmission end, particularly with poor HF conditions or aggressive choice of transmission speed.
There are two modem level mechanisms for determining end of transmission:

1. Modem Protocol. This is supported in STANAG 5069, but not older protocols. It definitively marks end of transmission.
2. EOM Marker. Two special bytes sent in the data stream.   Care needs to be taken to handle the case where the "real data" includes this value.   Heuristics to validate include ensuring that only "padding data" follows the EOM and that the RF signal falls off after the block is complete.

Both of these mechanisms can be "lost" due to fading or other data corruption.   A modem will detect loss of RF signal.   This can be an indication of transmission end or it could be a reception gap in a longer transmission.   A modem will generally wait for a period before considering the transmission ended due to loss of RF.   During this time, the modem will continue to send data to STANAG 5066 at "modem speed" and will maintain synchronization with the transmission (which may have ended).

The D_PDU EOT mechanism is helpful when the modem continues to receive in this way.   A STANAG 5066 server can determine the end of transmission time from any single D_PDU.   It will "know" that a transmission has finished, even when the modem continues to provide data (e.g., because Modem has not explicitly detected end of the RF transmission and is maintaining sync).   The DTS will be able to switch to transmission, while the modem is still sending it (spurious) data.

T.3.4. Modem Synchronizing During Transmission

HF Waveforms start with a robust synchronization sequence, and so normal expectation is that modems synchronize at start of transmission, so that data bytes are correctly aligned over the transmission, even when there is loss and corruption.

Modem protocols can resynchronize during a transmission.   This means that data can start to be received part way through a transmission.   Also, a transmission can be "lost" and then a latter part of the transmission picked up as a new transmission.

Some waveforms, notably STANAG 4539 and STANAG 4285, synchronize very well during when the initial synchronization fails and it is desirable to allow for this on reception.

STANAG 5069 does not synchronize well during a transmission, if the initial synchronization fails. This means that when STANAG 5069 is used, it may be desirable to use longer (and more robust) pre-ambles and **may** be desirable to avoid overly long transmissions.

T.4. Crypto Layer
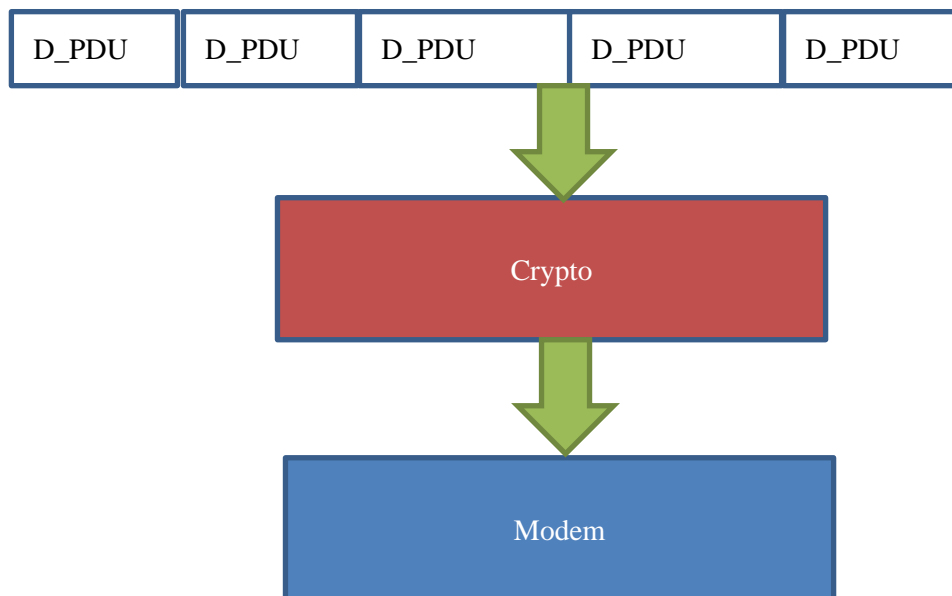
T.4.1. Service Interfaces

Figure T-3: STANAG 5066 Interface to Crypto

The basic service model of this protocol is that the data interface to and from the crypto layer is essentially the same and shown in Figure T-3.   This is in line with STANAG 5066 Annex D.

STANAG 5066 needs to be aware of the protocol overheads of the Crypto Layer, so that it can correctly calculate transmission lengths. It **may** also choose to align D_PDU boundaries to modem block boundaries, which can be done precisely with this protocol stack.

T.4.2. Crypto Mapping & Counter Mode

The Crypto needs to operate as a stream crypto.   Most modern encryption algorithms are block-oriented, which cannot be used directly.   This is addressed with Counter (CTR) Mode, which is a mechanism to provide stream encryption from a block cipher.   This is now widely recognized as a secure approach.

Counter mode works by initializing the cipher with an Initialization Vector (IV) and/or a Nonce. This then generates a crypto stream of bytes with as many bytes as needed. Sender and Receiver share the IV/Nonce and so can generate identical crypto streams.

The sender will XOR the data stream with the crypto stream to produce a transmission stream, which is sent between the modems. The receiver can then XOR the received transmission stream with the crypto stream to restore the data stream sent by the peer STANAG 5066 server. This received stream may have data corruption due to modem level HF errors.

The strict synchronization of modem data transfer ensures that the streams remain aligned over the complete transmission.

T.4.3. Crypto Initialization

This section considers crypto initialization in more detail.

T.4.3.1. General Model

There is information which needs to be shared between sender and receiver which is configured prior to the transmission. This might be an external mechanism such as pre-placed keys, or other external mechanism.

Then there is information provided for each transmission, which will typically be an IV and/or Nonce, but there could be other information for encryption mechanisms other than AES.

T.4.3.2. AES Initialization

Each transmission will provide a 2 byte reference to the AES Key/Nonce pair. This reference enables:

- Different keys may be used for different pairs of 1:1 communication and multicast/broadcast groups.
- New keys can easily be used in the event of key compromise.
- Giving keys a limited lifetime, with migration to new keys.

Management of this reference and distribution of AES Keys and Nonces is not specified in this annex.

Each transmission will include an 8 byte IV that is unique for the AES Key/Nonce pair. Uniqueness can be ensured by simple incrementing or by Linear Feedback Shift Register. The last IV used should be recorded on permanent storage, so the unique IV will be ensured in the event of restart.

T.4.3.3. Resilience over HF

Transferring the initialization information needs to address potential corruption on transfer. HF errors tend to cluster, so the best approach to resilience is to repeat the information at intervals in the data stream.

The protocol also needs to address the possibility of the initial modem synchronization not happening. This is done by use of extended information that includes:

1. A Maury-Styles two byte pair, that enables synchronization.
2. A counter to enable the amount of preceding data to be calculated. This will enable to full crypto-stream to be determined, so that the received data can be XOR'd from the correct point.

T.5. Crypto Layer Protocol

The Crypto protocol defines two PDUs.   The Crypto Sync PDU is specified in Figure T-4.

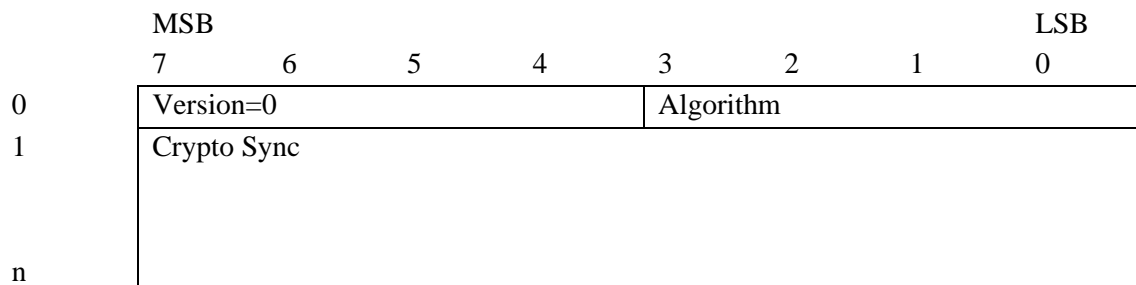|	MSB							LSB
	7	6	5	4	3	2	1	0
0	Version=0				Algorithm			
1	Crypto Sync							
n								

Figure T-4: Crypto Sync PDU

For this version of the protocol, Version=0.   The Algorithm identifies that Algorithm used. The length and semantics of the Crypto Sync bytes are determined by the Algorithm.

AES has Algorithm=0, and the encoding specified in Figure T-5 to give an 11 byte PDU.

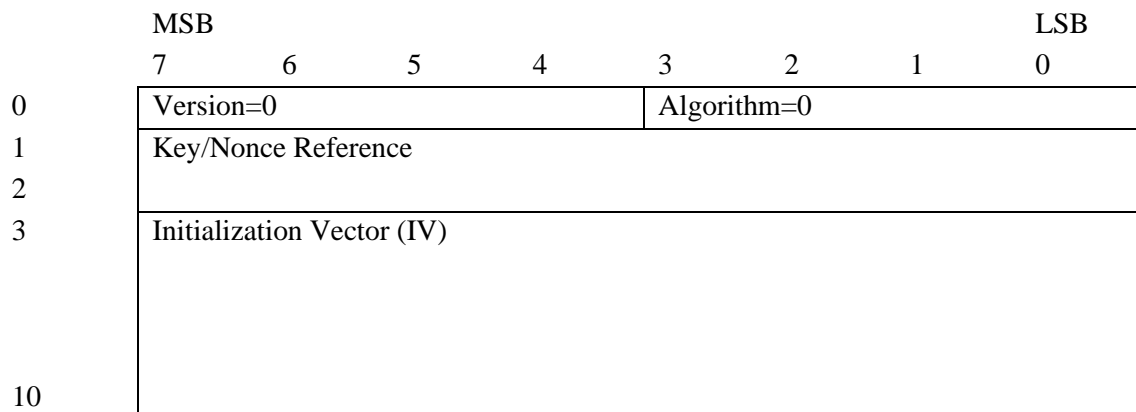|	MSB							LSB
	7	6	5	4	3	2	1	0
0	Version=0				Algorithm=0			
1	Key/Nonce Reference							
2								
3	Initialization Vector (IV)							
10								

Figure T-5: Crypto Sync PDU for AES

The encoding in Figure T-5 has the following elements:

1. The Key/Nonce Reference is two bytes and identifies the AES Key and Nonce pair to be used.
2. The Initialization Vector is an 8 byte IV.

The Extended Crypto Sync PDU is defined in Figure T-6:

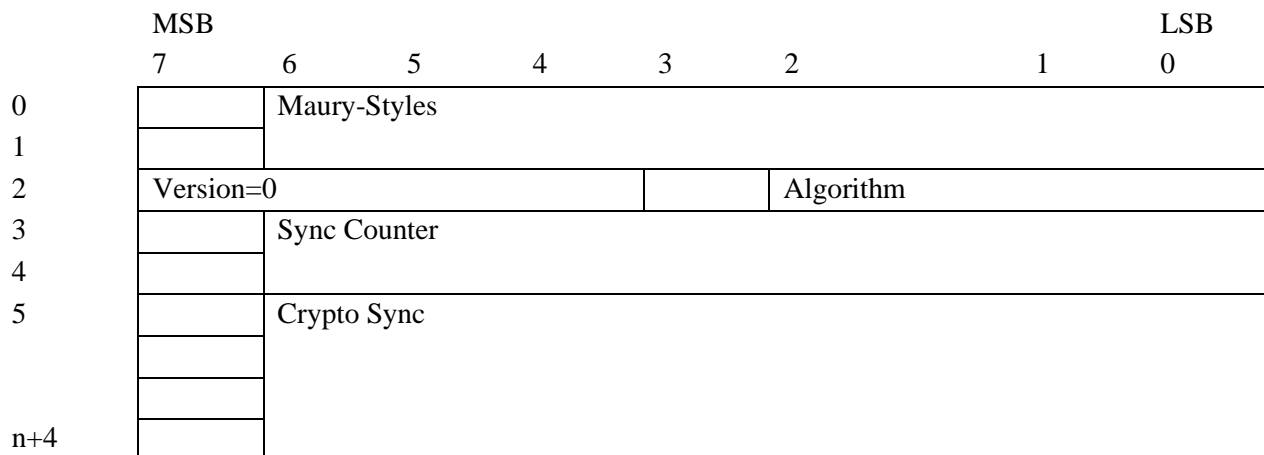|	MSB							LSB
	7	6	5	4	3	2	1	0
0		Maury-Styles						
1								
2	Version=0					Algorithm		
3		Sync Counter						
4								
5		Crypto Sync						
n+4								

Figure T-6: Extended Crypto Sync PDU

The encoding in in Figure T-6 has the following elements:

1. Maury-Styles is a two byte fixed header using the STANAG 5066 Maury-Styles value (90EB) using the same format as the one for D_DPUs defined in Annex C.

2. Sync Counter indicates the number of the Extended Crypto Sync PDU. The first one is numbered 1, the second 2, and so on. For AES, this PDU is 15 bytes long.

In order to protect against loss of Crypto Sync PDU, it is repeated in the transmission. Because transmission times can vary from less than a second to many minutes and speeds from 75bps to 240 kbps, there is a high variation of size of bounded stream. Repetition of Crypto Sync PDUs needs to be close enough to ensure repetition at the slowest speeds and to provide reasonable spacing at higher speeds to provide protection against corruption of full modem block.

Crypto Sync PDU	256 Bytes Encrypted Data	Crypto Sync PDU	256 Bytes Encrypted Data	Crypto Sync PDU	256 Bytes Encrypted Data

Figure T-7: Crypto Sync PDU repetition for first 4096 bytes of Data

For the first 4096 bytes of transmitted data, a Crypto Sync PDU shall inserted before every 256 bytes of

encrypted data, as shown in Figure T-7. If less than 4096 bytes of encrypted data is being transmitted, the last block **may** be less than 256 bytes. Figure T-7shows the start of a block transmitted data. For AES, the Crypto Sync in this region has an overhead of 4.3%.


Figure T-8: Extended Crypto Sync PDU repetition

If more that 4096 bytes of encrypted data is transferred, the initial 4096 bytes are transferred as shown in Figure T-7. After this, data is sent in 4096 byte blocks, with each block preceded by an Extended Crypto Sync PDU. The last block **may** be less than 256 bytes.

Figure T-8 shows the last 256 bytes of the first 4096 bytes of data, followed by two blocks of 4096 bytes. For this extended region, AES Crypto Sync overhead is 0.4%.

Annex U: IP Client

(Optional)

DRAFT 1.0 of 1st April 2020

This Annex specifies operation of IP (the Internet Protocol) over STANAG 5066, covering both IPv4 (RFC 791) and IPv6 (RFC 8200). This enables support of some IP applications over STANAG 5066.
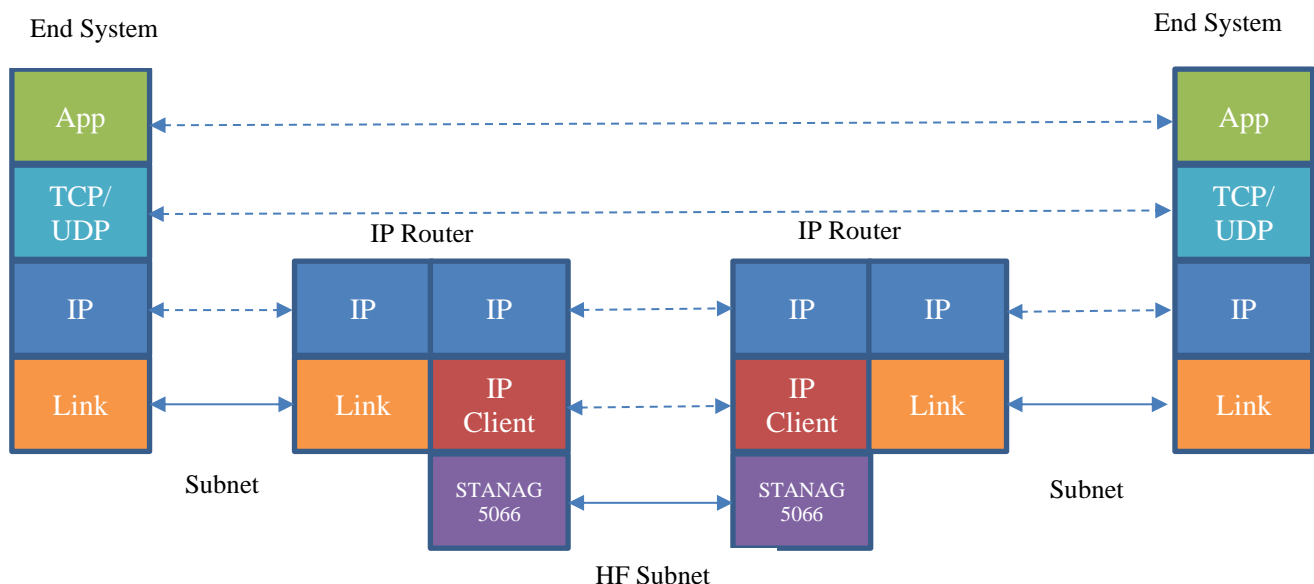
U.1.   Model of Operation

Figure U-1: IP Client HF Subnet Model

The general model of IP usage and IP Client is shown in Figure U-1. Applications communicate end to end, over end to end layer protocols such as TCP and UDP. IP is used to communicate over a series of one or more subnets. IP is a per-hop protocol with end to end implications. For any subnet, a link protocol appropriate to subnet technology is used to communicate over the subnet. IP Client, specified in this annex is used to communicate over an HF subnet using STANAG 5066. IP Client is a simple layer protocol to enable IP to be exchanged using STANAG 5066.

In most configurations, the IP Client will be logically associated with an IP Router, and not an application, end system or host.   It is possible to provide a router function with IP Client co-resident on an end system.

It is also possible to configure and end system to directly use IP Client, and for the end system to have an IP address on the subnet.

U.2.   Scope of Application

The model shown in Figure U-1 is general purpose, and in principle can be used to support any IP service running over an HF subnet using IP Client. In practice, the choice of IP service that can be usefully deployed over IP Client is limited. Under load, the interaction between TCP and an IP Client subnet is inefficient in many configurations, and for many deployments use of IP Client is not suitable for TCP and HTTP provision.

IP Client is suitable to provide a range of IP services, such as ICMP Ping, and some low volume protocols operating over UDP.

U.3.   General Requirements

An IP client implementation **shall** be capable of sending and receiving encapsulated IP datagrams with unicast (i.e., point-to-point) IP addresses, using both ARQ- and non-ARQ-transmission modes in STANAG 5066.

An IP client **may** be capable of sending and receiving encapsulated IP datagrams with multicast (i.e., point-to-multipoint) IP addresses, using non-ARQ-transmission modes.

U.4.   Encapsulation of the IP Datagram using STANAG 5066 Service

IP datagrams **shall**[1] be encapsulated within the U_DPU field of S_UNIDATA_REQUEST Service specified in Annex A, and delivered to clients at the destination node(s) within the U_DPU field of S_UNIDATA_INDICATION Service. There are no framing characters required or allowed.

The first byte of the header of the IP datagram **shall**[2] be aligned with the first byte of the U_PDU field within the primitive, and so on to the last byte of the IP datagram and U_PDU field.

The encoded bytes of an IP datagram submitted for transmission over the subnetwork **shall**[3] be bit-/byte-aligned with the octets in each U_PDU encapsulated in the S_UNIDATA_REQUEST primitive.

The least-significant bit (LSB) of each octet in the IP datagram **shall**[4] be aligned with the LSB of the U_PDU's octet.

IP_HDR (n bytes)	Byte[0]	-- IP_DATA (e.g., TCP, UDP) --	Byte[m-1]; m = k - n
Byte[0]		- - - U_PDU - - -	Byte[k-1]

Figure U-2: - Mapping of IP Datagram into an S_UNIDATA_REQUEST
U_PDU

U.4.1.   IPv4 and IPv6

Both IPv4 and IPv6 datagrams can be carried.   The version of IP used can be determined from the first four bits of the first byte of the IP Header.   Choice of IP version is transparent to this protocol, but IP Client implementations **may** need to control actions based on IP version. If needed, this version is straightforward to determine.

U.5.   IP-Client Subnetwork Service Requirements

IP clients **shall**[(1)] bind to the HF Subnetwork at SAP ID 9.

IP support by the HF subnetwork **shall**[(2)] be configurable to use either ARQ or non-ARQ delivery services.

Selection of the delivery mode (established in the S_UNIDATA_REQUEST message), and priority **shall**[(3)] be based on the type of IP address and other information on the IP datagram.

An IP client **may** set the subnetwork's default service requirements in the S_BIND_REQUEST message as a function of the most likely traffic that it expects to process.

An IP client **shall**[(4)] be capable of overriding the subnetwork's default Service Type requirements and dynamically setting the S_UNIDATA_REQUEST Delivery Mode for each IP Datagram submitted to the HF subnetwork in accordance with the requirements of Section F.12.4.

U.5.1.   IP Support using ARQ Service

HF subnetwork support using reliable point-to-point delivery between a pair of nodes is preferred for efficiency in the IP and higher-layer protocols, but in general cannot support IP-multicast protocols. [NB: The exceptional case is when an IP multicast address can be mapped to a STANAG 5066 unicast address, e.g., when tunnelling multicast traffic over an HF point-to-point link.]

The service definition for reliable-IP datagram delivery using the ARQ service **shall**[(1)] be as follows:

1. Transmission Mode =   ARQ,
2. Delivery Confirmation =   NODE DELIVERY,
3. Deliver in Order =   IN-ORDER   DELIVERY   or   AS-THEY-ARRIVE
4. Priority set according to Quality of Service as described in Section U.7.

U.5.2.　　IP Support using non-ARQ Service

If IP-multicast address groups are supported within the HF subnetwork environment (i.e., an application wishes to take advantage of the broadcast nature of the HF channel to support IP multicast), then the HF subnetwork **shall**[1] be configured in non-ARQ mode to support this requirement.

IP support using non-ARQ service **may** be used for IP unicast services.

For IP datagrams using non-ARQ service, the HF subnetwork service **shall**[2] be configured as follows:

1. Transmission Mode = non-ARQ,
2. Delivery Confirmation = none,
3. Deliver in Order = AS-THEY-ARRIVE.
4. Priority set according to Quality of Service as described in Section U.7.

The number of repeats for the D_DPUs in the service **may** be set to a value greater than one to provide some increased probability of receipt and reliability when using the subnetwork for IP multicast support.

U.5.3.　　Addressing

The IP Client Architecture, as shown in Figure U-1, interconnects IP subnetworks. This is anticipated to be the most common mode of deployment. This means that the IP addressing is at a higher level than STANAG 5066 addressing.   Two models of managing the address mapping are described.   An IP Client implementation **may** implement one or both modes.

An IP Client implementation will be associated with an IP router on the local subnet. Communication between IP Client and the IP router is an implementation choice.

U.5.3.1. Full Router Control

In this model, IP Client is configured to connect to a single peer.   The router will send IP packets to the IP Client, which will send the IP packet to its peer IP Client and associated subnet/router.   This enables the router to treat IP Client as providing access to the peer subnet.   This is a "link" in IP router terminology.   Routing can then be configured at the router level.   This can make use of routing protocols.   This approach is going to be sensible when a node can be used to route to many subnets.

U.5.3.2. Shared Router and IP Client Control

When IP Client on a node can directly address multiple peers, it **may** be convenient to have the router send it IP packets for all destinations rather than to have a logically separate communication for each link.   In this case, IP client can perform basic routing based on IP address subnet, in order to select the correct peer and associated STANAG 5066 address.   This is a convenient approach, provided that peer IP Clients either support a small number of fixed IP subnets or provide a default route (e.g., a shore station).

U.5.4.   Addressing with End System Connected to Subnet

It is possible to directly connect an end system to an HF Subnet. Where all nodes are configured in this manner, they will all have IP addresses allocated on the HF Subnet.   In this type of configuration, there is a direct 1:1 mapping between IP addresses and STANAG 5066 addresses.

U.5.5.   IP Datagram Queueing

A key choice for IP Client implementation is handling STANAG 5066 flow control, when the load of arriving IP packets is greater than the HF subnet can immediately handle.   A simple choice is to drop the IP packet when there is STANAG 5066 flow control. This is a simple approach, which follows the IP model of discarding packets on congestion and expecting the higher layers to adapt. Because of long HF delays, this discarding can lead to significant inefficiency, particularly when ARQ is used.

The alternate option is to queue arriving IP packets. This can lead to improved performance in some situations, but large queues building up with very long delays will lead to a different type of performance issues.   Handling this is a key implementation choice.

U.6.   Router Functions

There is a close relationship between IP Client and an associated router.   Consideration needs to be taken as to where some functions are provided.

U.6.1.   Segmentation and Reassembly

Any Unidata containing IP data **shall**[(1)] contain a complete IP datagram.

Segmentation and re-assembly, if used will generally be performed by the router.   Segmentation and re-assembly **may** be performed by IP Client.

U.6.2.   Internetwork Control Message Protocol

The IP Client and Router **shall** jointly ensure that various standard requirements are addressed between them:

1. That the Internet Path MTU discovery (PMTU) Protocol [RFC1191] for IP datagrams marked with the DON'T_FRAGMENT flag that also exceed the HF subnetwork MTU size.

2. That Internetwork Control Message Protocol (ICMP) is supported to provide MTU discovery.

U.7.   Selecting Options to Provide Best Quality of Service

For each IP packet handled, IP Client can made a number of choices:

1. ARQ vs Non-ARQ.

2. Priority.

3. Queueing strategy, as described in Section U.5.5.

4. Discard/Filtering.   IP Client may choose to discard certain IP packets.

Considerations for this choice can include general conditions, such as typical transfer speed and SNR. The most useful choices are based on each IP packet.   Options that **may** be considered include:

1. Destination Address.

2. IP Protocol (e.g., ICMP, GRE, UDP, TCP etc)

3. For IP protocols with Ports, in particular TCP and UDP, the port used.

It is **recommended** to have this choice configurable.   The following considerations may be helpful.

1. ARQ generally gives better performance than non-ARQ.

2. Some queuing generally leads to better performance, but very long queues need to be avoided.

3. Discarding traffic that is not explicitly allowed can be helpful.

4. Giving higher priority to control traffic such as ICMP is generally desirable.

U.8.   Changes in Edition 4

General descriptive text is added, including notes on recommended scope of application of IP Client.

This annex is changed from mandatory to optional.

Support for IPv6 is added. This needed clarification of procedure, but no change of protocol.

Support of multicast is made optional. This can be non-trivial to support, and its benefit is unclear.

Removal of controls related to Differentiated Services and TOS.   These are not relevant to modern IP networking and added unnecessary complexity.

Address handling rewritten, as Edition 3 text does not make practical sense. Addressing needs to be handled as subnet level.

Annex V: Compressed File Transfer Protocol
(Optional)

DRAFT 1.0 of 17th April 2020

Compressed File Transfer Protocol (CFTP) provides a protocol for transferring basic SMTP messages over STANAG 5066. It replaces the earlier HMTP protocol, which did not provide compression.

The Compressed File Transfer Protocol (CFTP) **may** be used to meet requirements in the NATO C3 Technical Architecture and the NATO Common Standards Profile (NCSP) for use of the SMTP Protocol. CFTP is used to reliably send compressed SMTP e-mail over a STANAG 5066 HF subnetwork from one message server to another.

CFTP is a vendor-defined standard derived from client definitions made in earlier, non-mandatory, versions of STANAG 5066 Annex F.

V.1.   General Requirements

Implementations of STANAG 5066 **may** provide a CFTP client. If provided, a CFTP client **shall**[(1)] conform to the requirements specified herein.

The CFTP protocol **shall not** be used for Formal or High Grade Military Messaging (i.a. military orders). For Formal or High Grade Military Messaging, ACP 142 as specified in Annex Q **shall**[(2)] be used. The CFTP protocol **may** be used for informal interpersonal e-mail only.

CFTP **shall**[(3)] operate within the node model shown, providing transfer services from one SMTP e-mail server to another via the CFTP client.
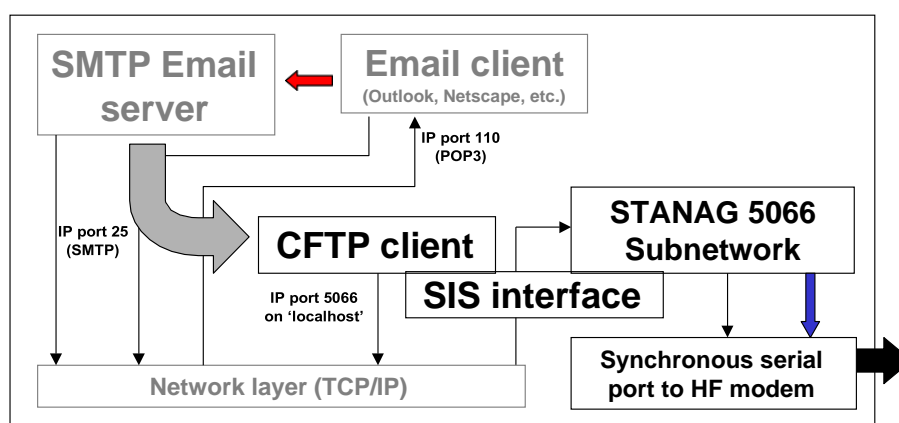

Figure V-1: CFTP Client and Node model

In general, the interaction of the CFTP client with the mail-server is beyond the scope of this STANAG. In operation, when an email message is received at a 5066 node, it is placed in an incoming mail folder (mail spool

directory). The CFTP client, also called the Delivery Agent (DA), removes mail from this incoming folder and processes the mail for delivery over HF via 5066. The CFTP DA compresses the message and information about the message, e.g. size, id, recipients, etc. into a file. This compressed file is then transferred to the destination 5066 node(s) using the original form of the Basic File Transfer Protocol (BFTP), referred to here as BFTPv1. The original BFTPv1 format is incorporated directly in the CFTP specification below to free it from dependencies on (and incompatibilities with) BFTP specification found in STANAG 5066 Ed3.

V.2.   CFTP Subnetwork Service Requirements

CFTP clients **shall** bind to the HF Subnetwork at SAP ID 12.

V.3.   CFTP Connection-Oriented Protocol

The CFTP application **shall** use the earlier form of the RCOP Protocol Data Unit ("RCOPv1") defined in the Figure below. [NB: As a historical note, this corresponds to the definitions of the RCOP protocol found in the original information-only Annex F Edition 1. It does not include the Application Identifier within the PDU Header that has been added in STANAG 5066 Edition 2.]
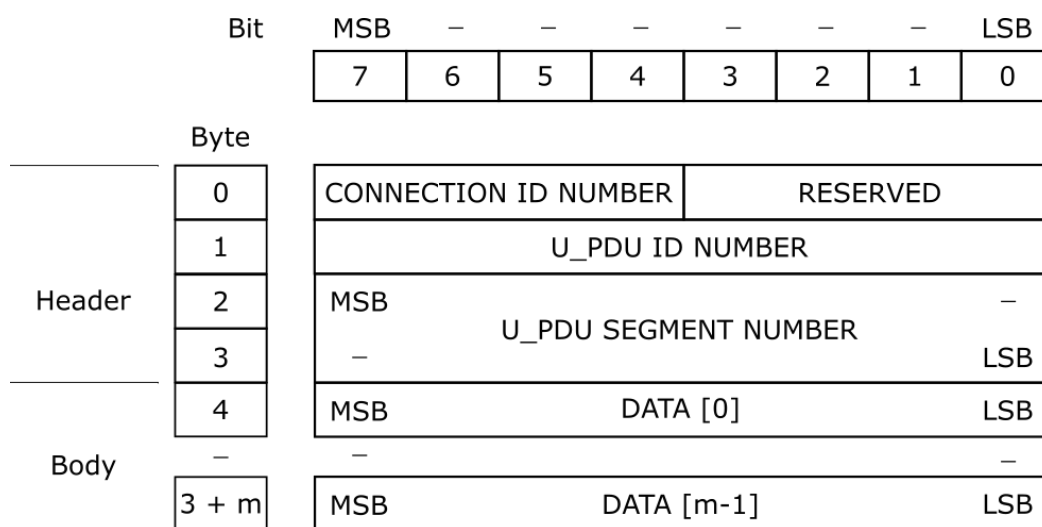

Figure V-2:. CFTP Protocol Data Units (identical in format to (Original) RCOP Protocol Data Units (RCOPv1) from STANAG 5066 Annex F Edition 1)

The following are required for RCOPv1 PDUs:

1. The connection ID number **shall** be a value from 0-15. Connection ID number 0 shall be reserved for non-multiplexed connections.
2. The reserved bits **shall** be set to 0.
3. The U_PDU ID numbers **shall** be assigned consecutively to U_PDU (i.e., files).
4. The U_PDU segment number **shall** be assigned consecutively to segments within a single U_PDU. The first segment transmitted **shall** be assigned segment number 0. If a U_PDU is not segmented, the single segment

transmitted **shall** be assigned number 0.

V.3.1.   Compressed File-Delivery and Delivery-Confirmation

Compressed files **shall**[(1)] be transferred from one CFTP client to another using the Edition 1 Basic File Transfer Protocol ('BFTPv1') as defined in the subsections below.

Client Delivery confirmation **shall**[(2)] be provided using the CFTP Message Acknowledgement, defined in the subsequent section (see Section V.3.3), as the body-part of the PDU.

In principle, up to 256 files could be transferred concurrently using the unique RCOPv1 U_PDU ID number for each transfer, using the U_PDU_ID as the identifier to match acknowledgements to the file acknowledged. As there is no negotiation protocol currently defined to determine if a given receiving node supports this capability, a sending node **must** have prior knowledge that a given receiving node supports concurrent multiple- file delivery.

Consequently, the Client Delivery confirmation protocol is nominally stop-and-wait — a new file **should not**[(1)] be sent with a given U_PDU ID until a message acknowledgement has been received. However, this recommendation **may**[(1)] be relaxed to allow concurrent multiple-file delivery when the sending node has prior

knowledge that the receiving node supports the capability. New implementations of CFTP **should** support concurrent multiple file delivery.

V.3.1.1.BFTPv1 Specification [NB: corresponding to the original Edition 1 BFTP specification]

The format for the basic-file-transfer-protocol data unit Version 1 (BFTPv1) **shall**[(1)] be in accordance with the following Figure, which defines a header part and a file-data part for the BFTP_PDUv1.

| BFTP_HDR | Byte[0] | - - - FILE_DATA[] - - - | Byte[p-1] |

Figure V-3:: Basic FTP Version 1 Protocol Data Unit (BFTPv1 PDU)

The detailed structure of the BFTPv1_PDU **shall**[(2)] be in accordance with the following Figure, and provide the following information fields:

1. BFTPv1_PDU Header Part:
   - SYNCHRONIZATION - two bytes corresponding to the control bytes DLE (Data Link Escape) and STX (Start of Text).
   - SIZE_OF_FILENAME - one octet in size.
   - FILE_NAME - a variable length field, equal in size to the value specified by the SIZE_OF_FILENAME field.
   - SIZE_OF_FILE - a four-octet field.

2. BFTPv1_PDU Body Part:
   - FILE_DATA[] - a variable length field, equal in size to the value specified by the
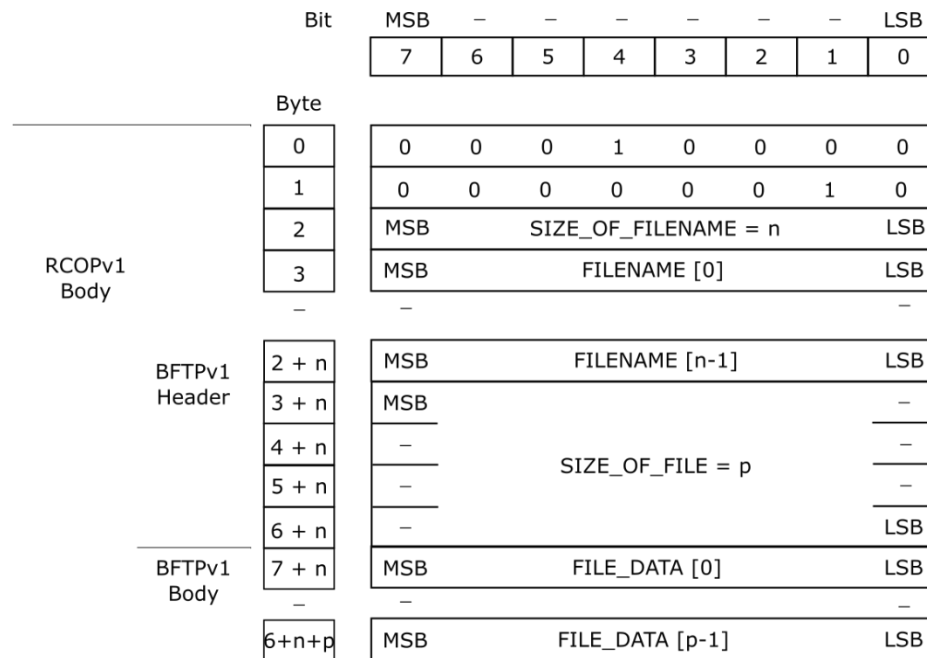
SIZE_OF_FILE field.

Figure V-4: BFTPv1 Protocol Data Unit Structure

The SIZE_OF_FILENAME field **shall**[(1)] be a 1-octet fixed-length field whose value (n) **shall**[(2)] equal the number of octets used to encode the FILENAME field.

The FILENAME field is **shall**[(1)] be a variable-length field, the size of which **shall**[(2)] be specified by the value (n) of the field SIZE_OF_FILENAME. This field represents the name of the file sent using the Basic File Transfer Protocol. The first byte of the filename **shall**[(3)] be placed in the first byte of this field, with the remaining bytes placed in order. The semantics of file names and naming conventions are beyond the scope of this STANAG (e.g., there is no requirement that the filename be a null-terminated character string.)

The SIZE_OF_FILE field **shall**[(1)] be a 4-octet fixed-length field whose value **shall**[(2)] specify the size (p) in octets of the file to be sent. The first octet of the SIZE_OF_FILE field **shall**[(3)] be the highest order byte and the last byte the lowest order byte of the field's binary value.

V.3.2.   BFTPv1 Segmentation and Reassembly Requirements

If the BFTPv1_PDU exceeds the maximum size of the data field permitted in the RCOPv1 PDU (i.e, if the CTFP_PDU is larger than the MTU_size less 4 octets (i.e., MTU-4)), the CFTP client **shall** segment the BFTPv1 PDU, placing successive segments in RCOPv1 PDUs (original Edition 1 format) with consecutive U_PDU sequence numbers.

When received, the CFTP client **shall**[(2)] reassemble the BFTPv1 PDU if it determines that the BFTPv1 PDU has been segmented. Subject to local-host file naming conventions, the CFTP client **shall**[(2)] store the received file with

the name transmitted in the header with the file. [*NB: there is no guarantee therefore that the file will be stored on the destination host with the same name that it was sent.*]

V.3.3. Message Acknowledgement

Client Delivery confirmation **shall**[1] be provided using the Message Acknowledgement defined below, sent as the body-part of a CFTP/RCOPv1 PDU.

	Bit	MSB	-	-	-	-	-	-	LSB
		7	6	5	4	3	2	1	0
	Byte								
BFTPv1 ACK	0	0	0	0	1	0	0	0	0
(RCOPv1 body part)	1	0	0	0	0	1	0	1	1

Figure V-5: BFTPv1 Message Acknowledgement Structure

On receiving the last byte of the message, the receiving client **shall**[2] send the Message Acknowledgement (0x10 0x0B) — with the same RCOPv1 U_PDU_ID NUMBER and CONNECTION ID NUMBER as the message being acknowledged — to confirm that the entire message has been received. (N.B. This is equivalent to the "ZEOF" message of the Z-modem protocol.)

V.4. CFTP Compression/Decompression

The compressed file **shall** be created and decompressed in accordance with RFCs 1950, 1951 and 1952 (i.e., the gzip utility defined in RFC1952).

V.5. CFTP Compressed File Data Format

The compressed file data **shall** be formatted as a series of fields. The fields are separated by the linefeed <LF> character 0x0A. The fields and the order in which they are compressed are described in the following table.

Table V-1 CFTP Mail File Structure

Order of Compression	Field Name	Description
1	MessageID	The MessageID field is represented by an arbitrary string that serves as the ID for the message. The MessageID is unique to an e-mail message. It is not the same as the ID in the email message that follows "Message-ID:" in the header. When the compressed file is decompressed, the MessageID is used as the root filename for the decompressed components. The MessageID must be less than 256 characters and is composed of upper/lowercase alphanumeric characters.

2	RecipientList	The RecipientList is a string containing e-mail addresses extracted from the e-mail message, each address delimited by the "," character (0x2c). The first address in the recipients list is the "Return-Path". There can be cases where there is no return path, e.g. the mail is being bounced by a Mail Transfer Agent. In these cases, the first address will be an empty string (i.e., either a single space [0x20] character or no characters at all) and it will be followed by a comma (0x2c). The recipients list must be less than 10240 characters.
3	MessageSize	The MessageSize is encoded as a decimal number in string format. It represents the size (in bytes) of the Message field that follows the MessageSize field.
4	Message	Actual message as received by an SMTP receiver *i.e. including the terminating sequence <CRLF>.<CRLF> and any additional characters that may be required for transparency as defined by RFC 821 para 4.5.2* .

Note 1. All characters are 8 bits.
Note 2. The terminating sequence **<CRLF>.<CRLF>** is that shown in Example 1 of RFC 821 and equates to the 5 ASCII Characters with codes, in hexadecimal, of 0x0D, 0x0A, 0x2E, 0x0D, 0x0A.

V.6.　　Detailed Description of CFTP

1) An e-mail client is used to send an e-mail to an SMTP server.

2) The CFTP application extracts the e-mail message from the directory in which it was placed by the SMTP server. An example e-mail message in the correct format is shown in Figure F-23 below.

3) The CFTP mail-file shall be built as follows:

<MessageID><LF>　　　　//　The MessageID field **shall**[(1)] be represented by a character string. The MessageID **shall**[(2)] be unique to an e-mail message. It is not the same as the ID in the email message that follows "Message-ID:" in the header. When the compressed file is decompressed, the MessageID **shall** be used as the root filename for the decompressed components. The MessageID **must** be less than 256 characters and **may** be composed of upper/lowercase alphanumeric characters.

<RecipientsList><LF>　　　　//　The RecipientList **shall**[(1)] be a character string containing e-mail addresses extracted from the e-mail message, each address separated by "," character (0x2c). The first address in the recipients list **shall**[(2)] be the "Return-Path". There **may** be cases where there is no return path, e.g. the mail is being bounced by a Mail Transfer Agent. In these cases, the first address **shall**[(3)] be an empty string (i.e., either a single space [0x20] character or no characters at all) followed by a comma (i.e, a "," character with octet value = 0x2c). The recipients list **must** be less than 10240 characters.

<MessageSize><LF>　　// The MessageSize **shall** be encoded as a decimal number in string format terminated by the linefeed character. It represents the size (in bytes) of the Message field that follows

<Message>　　　　// The Message field **shall** contain the e-mail message body part(s) extracted from the SMTP envelope.

4) The CFTP message (including header) shall be compressed in accordance with RFCs 1950, 1951 and 1952 using an application such as gzip.

5) The compressed CFTP message shall be encapsulated within a BFTPv1 PDU (i.e., it has a BFTPv1 header prepended to it, and the CFTP message shall be byte aligned within the FILE_DATA[] field of the BFTPv1 PDU.

6) The BFTPv1 message (i.e., BFTPv1 PDU) shall be segmented if necessary.

7) Each BFTPv1 PDU segment shall have an RCOPv1 header added (in accordance with Annex F.14.3.).

8) Each RCOPv1 packet shall be packaged into an S_UNIDATA_REQUEST and transferred using a Soft Link Data Exchange.

9) On reception the BFTPv1 message shall be reassembled, if required, and decompressed using a method compliant with RFC 1952 and the CFTP message reconstructed.

10) The received email messages shall be forwarded to an SMTP server using a standard SMTP dialogue based on information extracted from the CFTP header and inserting the "message" field into the payload of the SMTP message generated.

Received: from northampton (unverified [127.0.0.1]) by northampton.pdw<CRLF>
 (Rockliffe   SMTPRA   4.2.4)   with   SMTP   id   <B0000000133@northampton.pdw>   for
<root@essex.pdw>;<CRLF>
 Wed, 9 May 2001 12:09:03 +0100<CRLF>
Message-ID: <001f01c0d878$74da90d0$0d02a8c0@pdw><CRLF>
From: "Northampton" <administrator@northampton.two><CRLF>
To: <root@essex.pdw><CRLF>
Subject: Test <CRLF>
Date: Wed, 9 May 2001 12:09:03
+0100<CRLF> **MIME-Version: 1.0**<CRLF>
Content-Type: text/plain;<CRLF>
       charset="iso-8859-
       1"<CRLF>
Content-Transfer-Encoding:
7bit<CRLF> **X-Priority: 3**<CRLF>
X-MSMail-Priority: Normal<CRLF>
X-Mailer: Microsoft Outlook Express 5.50.4522.1200<CRLF>
X-MimeOLE: Produced By Microsoft MimeOLE V5.50.4522.1200<CRLF>
<CRLF>
This is the body of the test email<CRLF>
<CRLF>
.<CRLF>

Figure V-6: Example email in the Correct Format

The red and blue text above is the message body with text in blue being the terminating sequence <CRLF>.<CRLF> i.e 0x0D, 0x0A, 0x2E, 0x0D, 0x0A.