AMCP WG-M2 Appendix K

AERONAUTICAL MOBILE COMMUNICATIONS PANEL (AMCP)

Working Group M
2 nd Meeting
Järvenpää, Finland

25 April-3 May, 2001

Report of activities of subgroup dealing with VDL Mode 4 ATN Issues

AERONAUTICAL MOBILE COMMUNICATIONS PANEL (AMCP)

Working Group M
2nd Meeting
Järvenpää, Finland
25 April-3 May, 2001

Agenda Item 11: Review of results of subgroup dealing with VDL Mode 4 ATN issues

Report of activities of subgroup dealing with VDL Mode 4 ATN issues

Prepared by N. Fistas, Eurocontrol

SUMMARY

This Working Paper is a report on the activities of the subgroup that was established to examine the ATN part of VDL mode 4. As a result of these discussions, three amendment proposals have been submitted and proposed to WGM for approval. For some issues, the discussion is ongoing and the status is reported to WGM.

1 Introduction

1.1 Background

The 7th meeting of the AMC Panel approved SARPs for the VDL Mode 4 system “to support ATN compliant subnetwork services for surveillance applications”. The VDL Mode 4 system was validated in ICAO for surveillance applications (ADS-B and ADS-C). However, the VDL Mode 4 system is capable to support other applications and in particular point to point communications (i.e. CPDLC).

AMCP/WGC in its first meeting (October 2000) discussed the additional capabilities of VDL Mode 4 and has recognised the potential for VDL Mode 4 to support time-critical communication services as an ATN sub network as well as broadcast communications. Therefore, AMCP/WGM was invited to identify the activities and steps necessary to validate this position and to identify the appropriate actions to include communication functions within the VDL Mode 4 SARPs.

AMCP/WGM decided in its first meeting (December 2000) to establish a subgroup to examine further the issue. In particular the examination will address two aspects. The first one is the resolution of a number of issues as identified in Appendix O of the report of the 1st meeting of WGM. The second one is a comparative analysis of the performance of the different candidate systems to support point to point communications.

1.2 Scope

This report considers the first aspect of the above discussion: the optimisation of some VDL Mode 4 features and the provision of some additional functionalities to support a more efficient operation for point to point ATN applications.

This report presents the outcome of the subgroup discussions.

Based on the issues described in Appendix O of the report of the 1st meeting of WGM and the discussions in the same meeting, the subgroup set out to examine the following issues:

- the use of the Frame Mode SNDCF (and the associated issues: i.e. generation of join and leave events)
- concerns in association with the possibility of using a non unique address in ATS communications
- definition of VDL Mode 4 profiles describing the usage of particular options of 8208

2 Meetings, Participants and References

There were 3 teleconferences organised to discuss the issues:

- T1, 26 January 2001
- T2, 22 February 2001, and
- T3, 12 March 2001.

The participants in these meetings were:

- Eurocontrol (N. Fistas, M. Shorthose, T. Whyman, N. Wit)
- LFV (L. Johnsson, H. Malen)
- DFS (A. Schlereth)
- FAA (J. Hamelink, B. Philips)
- ADSI (P. Nair, S. Friedman, S. Heppe)
- CNS systems (T. Bergstrom, C. Nilsson, F. Qvigstad)
- SaabTranspondertech (C. Ericsson)

The following reference documents were produced and discussed in the frame of the teleconferences:

[1]: T. Whyman, “Using VDL Mode 4 as an ATN Subnetwork”, Eurocontrol. (Attachment 1)

[2]: M. Shorthose, “Options on the way forward for ATN interface in VDL Mode 4”, Eurocontrol. (Attachment 2)

[3]: M. Shorthose and N. Fistas, “FEC modelling in VDL Mode 4”, Eurocontrol. (Attachment 3)

[4]: S. Friedman, “FEC selection issues”, ADSI. (Attachment 4)

[5]: S. Friedman, “Change Proposal: FEC”, ADSI. (Attachment 5)

[6]: F. Qvigstad, “FEC in VDL Mode 4”, Sectra/CNS systems. (Attachment 6)

A final discussion on these issues took place in the meeting of the VDL Mode 4 subgroup during the WGM/2 meeting (WGM2/VM4). This paper presents the final status of the discussions.

3 Resolution of issues

3.1 Use by VDL Mode 4 of the ATN Frame Mode SNDCF

In preparation for the discussions, the use of Frame Mode SNDCF in VDL Mode 4 was analysed in [1] and the options were described in a briefing paper [2].

In the first teleconference T1, based on the above documents, the group agreed on the benefit of introducing the Frame Mode SNDCF in VDL Mode 4. The group discussed whether it would be useful to maintain the 8208 approach and it came to the conclusion that there is no benefit in keeping both approaches especially since today there exist problems in this approach in VDL Mode 2.

A draft Amendment Proposal to the VDL Mode 4 Technical Manual in order to introduce the Frame Mode SNDCF in the place of 8208 was discussed during the WGM2/VM4 meeting and a final proposal was agreed. This is provided in Working Paper 19 (WGM2/WP19).

The analysis in [1] highlighted two additional issues in relation to the ATN communications in VDL Mode 4:

- performance enhancements by using an FEC scheme, and
- potential benefits by using an alternative DLS protocol instead of the AVLC.

Although these issues are not in Appendix O of the WGM1 report, the group considered them as described below.

3.1.1 FEC and VDL mode 4

The FEC issue was discussed at length in all three meetings. In an ADS-B context, the VM4 VSG had decided in the past, that there is no need for an FEC scheme due to the periodic repetition of the information, the short length of transmissions and the capacity penalty (bits to accommodate FEC). Therefore this group considered the use of FEC only in the context of point to point communications. The FEC would be beneficial particularly for lengthy messages and for environments with increased traffic, in which the penalty of retransmissions could be significant.

In order to evaluate the benefit of using an FEC scheme, a simplified generic model [3] was produced and reviewed by the group. Some participants were of the opinion that performance enhancements similar to what an FEC scheme would provide, could also be achieved by specific implementation choices. However these choices are not covered in the SARPs and therefore they do not represent the baseline system performance.

An inclusion of FEC scheme has many implications that need to be thoroughly examined. The potential benefits of an FEC scheme versus the difficulties of implementing it were discussed at length. In the second and third meeting the modelling was reviewed and a number of parameters (i.e. BER and message length) used in the model were agreed. The group finally agreed in the third teleconference that there is benefit to include an FEC scheme to cover the ATN communications in VDL mode 4.

However there were diverging views on how this could be achieved. Reference [4] provided useful input in the discussions and highlighted the complexity of the problem. It became clear that the group could not resolve the issue in a teleconference meeting. Nevertheless, in order to advance further the discussion, group members committed to provide proposals for discussion in a proper meeting in the frame if the WGM/2 meeting. Furthermore, the group thought that the issue is of great importance and with many implications and therefore the advice of WGM should be sought. The draft proposals in relation to an FEC scheme are provided in [5] and [6].

During the WGM2/VM4 meeting the issue was again discussed. Overall there were three possible paths identified:

- redesign of the physical layer to accommodate FEC (Attachment 6 which is an example for the first option was presented)
- introduction of FEC for some transmissions (differentiated by a different synchronisation sequence
- introduction of FEC in some upper layers

There was no consensus in the way forward among the participants.

3.1.2 Definition of a simpler DLS protocol

The second additional issue that the group considered is the optimisation of the DLS. The group examined two main options as described in [2]. One is to maintain the AVLC protocol, which is specified in the VDL Mode 2 SARPs. However, [1] highlighted some drawbacks and the need for editorial work to make the specification “cleaner” and easier to read, understand and use in VDL Mode 4. The reason for the latter is that currently the use of AVLC in the VDL Mode 4 documents refers to the VDL Mode 2 documents together with a list of differences/exceptions. A draft document with the proposed changes have been developed and this editorial only aspect of the work resulted in a significant amount of text (~30 pages) to be introduced in the VM4 TM (text from the VDL Mode 2 TM). The other option is to define a simpler DLS protocol, which will only provide the required functionalities and avoid duplication. As a basis for this simpler protocol, the DLS protocol of VDL Mode 3 was used and in a separate working paper (WGM2/WP18) the status of this activity is reported.

During the WGM2/VM4 meeting, Working Paper 18 was presented and discussed. The group agreed that the new DLS is much simpler than the AVLC protocol. However, the group decided not to propose a new DLS protocol, until the analysis and testing of this protocol is complete.

The group agreed to clarify an issue that was identified in [1] (responsibility of retransmissions). For this an Amendment Proposal was discussed and agreed by the group. This proposal is submitted in a separate working paper (WGM2/WP32).

Even if the group decided to keep the old DLS protocol (AVLC), the group agreed that it would be beneficial to consider the new DLS protocol for future use. A question that remains is the feasibility/practicality of introducing a different DLS protocol after the publication of the Technical manual.

3.2 Use of non-unique address

The VDL M4 specification allows the use a non unique station address (Table 1.55 of the Technical Manual). This and the conditions of usage is described in Section 1.4.2.1.2 of the Technical Manual.
This may be a useful feature in an ADS-B environment. In the TM there is a note (Note 2) that states “It is anticipated that this (feature) will be limited to aircraft not receiving or requesting an ATC service”.

The group discussed the issue in T1 and decided that that this note is not sufficient if VDL mode 4 were to be used for ATN communications. The group proposes that that this note should become a “shall” to enforce the correct use of address in an ATC environment.

The issue is further analysed in Working Paper 15. In addition the issue was discussed again at the WGM2/VM4. Text for this amendment was approved in the meeting. In addition an additional change was approved, which provides an easy way to distinguish between aircraft and ground vehicles. Since both these changes are in the same section they are presented together in one amendment proposal in Working Paper 31.

3.3 ISO 8208 profiles for VDL Mode 4

This is an issue because the 8208 allows for optional implementation of certain features. This was identified in the VDL Mode 2 implementation activities and is an issue that affects the interoperability of different implementations.

J. Hamelink took the action in T1 to propose such a profile. A profile was proposed in a separate working paper (WP16). However, in the discussion of the Frame Mode SNDCF (see section 3.1), the issue of removing completely the 8208 interface was raised. Finally, in the WGM2/VM4 meeting, the use of Frame mode SNDCF was approved and the use of 8208 was removed from the VM4 Technical Manual. Therefore there is no need for such profiles for VDL Mode 4.

4 Conclusions

A subgroup was established to examine a number of issues identified in the 1st meeting of WGM in relation to the use of VDL Mode 4 for ATN ATS communications.

The discussion in the group resulted in proposing three amendment proposals (introduction of Frame Mode SNDCF, conditions of use of non-unique address, clarification in AVLC). The group discussed two other issues (use of an FEC scheme, and use of a new DLS protocol) for which there was no conclusion.

5 Recommendations

It is recommended that WGM:

- takes note of the analysis of the considered issues
- approves the three amendment proposals as described in the working papers 19, 31, and 32.
- provides advice/comments in relation to the potential use of an FEC scheme and of a new DLS protocol.

EUROCONTROL

ATN PROJECT

Using VDL Mode 4 as an ATN Subnetwork

ATN Ref. : DIS/COM/ProATN_Sup/DCI/AW_90

Author : Tony Whyman

Rev. No. : Issue 1.1

Date : 03/01/01

DOCUMENT CONTROL LOG

SECTION	DATE	REV. NO.	REASON FOR CHANGE OR REFERENCE TO CHANGE
	03/01/01	Issue 1.1	

TABLE OF CONTENTS

1. Introduction

1.1 Background

VDL Mode 4 is a class of subnetwork defined by ICAO to operate in the VHF Aeronautical Band. It has been developed in support of Surveillance Applications (ADS-B) but also supports a data transfer capability for both air to air and air-ground communications.

ICAO policy has been to provide a common Air/Ground data network access service according to ISO/IEC 8208 (a.k.a. ITU-T Recommendation X.25). However, this has proved costly in terms of both communications overhead and software development and has caused particular problems in VDL Mode 2 with respect to Ground Station Handoff, where the data link Handoff is exposed to the ATN Network Layer. In response, the ATNP has developed a new strategy for using Air/Ground data networks as part of the Aeronautical Telecommunication Network (ATN), which avoids the need for an X.25 network access service.

This new strategy is known as the “Frame Mode SNDCF” and permits direct use of air/ground datalink services without the need for a network access service.

1.2 Scope

This paper reviews the current state of the VDL Mode 4 SARPs and discusses how it may be used as part of the ATN via the Frame Mode SNDCF without having to develop and X.25 network access service.

1.3 References

1. ATNP/WG2/WP597 Proposed SARPs for a frame mode SNDCF

2. AMCP/7 – WP/81 Proposed VDL Mode 4 SARPS and Technical Manual

1.4 Document Structure

Chapter One – this chapter provides an introduction to the document.

Chapter two – Summary provides a summary of the documentation and the conclusions of the analysis.

Chapter three – Analysis of VDL Mode 4 SARPs provides a review of VDL Mode 4 highlighting issues with respect to ATN integration.

Chapter four – The Frame Mode SNDCF provides an overview of the Frame Mode SNDCF.

Chapter five – Accessing VDL Mode 4 Using the Frame Mode SNDCF provides discussion and proposals on the options for using VDL Mode 4 with the Frame Mode SNDCF.

2. Summary and Conclusions

The primary focus of this paper has been to identify how the ATN can be interfaced to a VDL Mode 4 network. However, part of this work has also been the analysis of VDL Mode 4 and how it operates as a data network.

The paper thus starts by analysing VDL Mode 4 and a number of issues have resulted from this analysis, which are recorded below. It also presents the Frame Mode SNDCF, as defined by the ATNP, and then shows how VDL Mode 4, as described by the current VDL Mode 4 SARPs can be used as an ATN subnetwork, and identifies a number of issues to be resolved.

2.1 VDL Mode 4 as an ATN Subnetwork

The most basic conclusion of this work is that VDL Mode 4 as defined by the ICAO Technical Manual can provide the functionality needed to support the ATN using the new Frame Mode SNDCF. No additional specification is required in terms of the service provided. However, a number of issues have been identified with respect to VDL Mode 4 offering a datalink service suitable for ATC use:

1. Whether the Join Event is to precede or follow Data Link Establishment. The issue here is about the maximum practical frame size for an XID frame and making the Join Event follow Data Link Establishment is a means for minimising the XID frame size by not concatenating it with Frame Mode SNDCF information. However, concatenation is more data link efficient and is otherwise desirable. The underlying issue here is about the lack of Forward Error Correction (see 4 below).

2. The need to fix the VDL Mode 2 inherited problem with regard to the timely generation of Leave Events on the ground. VDL Mode 4 can inherently provide a better solution to the Leave Event Latency problem than VDL Mode 2 by utilising ADS-B information (or the lack of such information) to identify when an aircraft has flown out of range. However, a VDL Mode 4 SARPs change will be required to do this.

3. The lack of prioritised data transfer. At present VDL Mode 4 does not implement priority between stations. This will limit its utility as an Air/Ground network to non-time critical applications and make it difficult to utilise the full capacity of each channel. A similar limitation is also true of VDL Mode 2.

4. There are issues surrounding whether the AVLC is the most appropriate datalink layer service for VDL Mode 4 which affects the way VDL Mode 4 is used as part of the ATN. At present, the need for short frames implied by the lack of Forward Error Correction (FEC) means that the only mechanism for improving performance is the AVLC provided capability to concatenate short frames together into a single transmission, and to force selective retransmission of any so concatenated frames received in error.

   However, if FEC was available then this would probably not be necessary especially as concatenation is also a feature of the Frame Mode SNDCF. This, in turn, may permit a simpler data link layer providing no more than a CRC and an alternating flag for detection of re-transmissions and hence rejection of duplicates.

2.2 VDL Mode 4 Specific Issues

In reviewing the VDL Mode 4 specification a number of issues have been identified with respect to the use of VDL Mode 4 as a data network in support of Air/Ground communications:

1. The lack of a Forward Error Correction (FEC) mechanism is a major weakness affecting the performance and usability of the network. It is understood that an appropriate expected error rate for VDL Mode 4 is 1 in 10^4 bits transmitted. This figure has been predicted for normal operations at an altitude where multipath effects will be limited and assumes a signal to noise ratio of the order of 12db.

   The lack of Forward Error Correction implies that, for example, the probability of successfully transmitting a 1KB message without error is approaching zero. The packet size has to be limited to small packets with 128 bytes probably being the realistic maximum – although the probability of successful delivery is then still less than 90%.

   The predicted error rate could easily rise if Co-Channel Interference due to simultaneous transmission by apparently CCI immune transmitters is greater than expected. This could be due to local effects or other reasons such as aircraft banking during transmission or reception.

   Performance is thus believed to be an issue and the provision of an FEC for use by the datalink service should be reviewed. Without an FEC, it is questionable as to whether the performance of the network will be sufficient for ATC use.

2. The AVLC Specification needs to be developed further. At present, too much is simply inherited from VDL Mode 2 without qualification. For example:

   - the VDL Mode 2 problem with respect to the “Leave” Event latency has been imported rather than fixed, as discussed above. Under VDL Mode 4, the periodic broadcast provides a simple mechanism by which a Ground Station can detect that an aircraft has gone out of range. This should be specified as the means to provide a “Leave” Event.

   - It is unclear whether the MAC Layer or the AVLC is responsible for retransmission. Resolution of this ambiguity in favour of the MAC layer (which it should be to take maximum advantage of VDL Mode 4) will result in substantive changes to the underlying HDLC procedures.

   It should be recognised that the VDL Mode 4 AVLC is not the same as the Mode 2 AVLC and “specification by exception” leads to confusion and probably error. The AVLC for VDL Mode 4 should be fully specified in the VDL Mode 4 SARPs.

3. While VDL Mode 4 does specify priority based queuing of packets for transmission, it does not provide a mechanism for prioritising access to the channel by competing stations. This implies that it is not going to be readily possible to offer performance guarantees to high priority services while simultaneously using the channel for lower priority applications. In consequence, channel utilisation and overall performance may well be inferior to VDL Mode 2 given that the actual data rate is slower.

4. Slot Management is an issue that needs to be addressed in order to provide prioritised access to the medium. At present, the Air/Ground Datalink Service specifies symmetric procedures for both airborne and ground systems. However, priority can almost certainly be better managed with a greater bias towards ground control. Furthermore, the current specification appears to require that all airborne systems are equally called upon to ensure proper operation of the service. This limits the potential to evolve and improve the slot management algorithms as each change will require new SARPs and upgrade of aircraft systems. If slot management is largely ground driven then the detailed algorithms for slot allocation (i.e. deciding which station gets which slot) do not have to be detailed in SARPs and improvement can be readily deployed by changes to a Ground Station only.

5. Many of the above issues appear to be due to the competing interests of surveillance and data communications. It understood that it is anyway expected that separate

channels will be required for ADS-B and data communications, and it may be that the specification should evolve differently for ADS-B and data communications purposes.

2.3 Conclusion

The existing VDL Mode 4 specification should be usable as an ATN subnetwork, at least for trials purposes. However, it is believed that the VDL Mode 4 specification does need further development before operational use can be considered.

3. Analysis of VDL Mode 4

3.1 Overview

VDL Mode 4 (VDL4) is specified for operation in the existing VHF Aeronautical waveband using the current 25kHz channel spacing. Digital transmissions are modulated using Gaussian-filtered Frequency Shift Keying (GFSK) with a data rate of 19.2kbits/s.

Under VDL4 each communications channel is a shared transmission medium with both airborne and ground based users. Media Access procedures are based on a time slotted access control scheme, with synchronisation provided by a UTC time source in normal operation. There is no controlling station although a Ground Station can “autotune” an aircraft to a different frequency. The principles behind the access control mechanism can be traced back to Slotted Aloha principles. However, it is much improved from such a simple scheme by the use of slot reservation protocols and CSMA techniques.

The MAC layer service provides direct support for an ABS-B application and a data link level service derived from the VDL Mode 2 AVLC.

Multi-channel operations is also part of the specification – although this will require airborne radios with typically three receive and one transmit channels.

3.2 VDL4 Principles of Operation

The unique feature of VDL4 is the MAC scheme used to provide for shared access to the communications channel. This is based on based on a combination of :

a) A time slotted p-persistent access control scheme to provide an underlying structure for co-ordinating transmissions, combined with “Listen before send” when transmitting into an unreserved slot (Random Access procedures)

b) a combination of schemes used to reserve access to transmission slots and hence to perform collision avoidance and improve upon the performance of the network.

3.2.1 Channel Structure

The *SuperFrame* provides the fundamental channel structuring concept. A SuperFrame comprises an integral number of transmission slots transmitted during a period of 60 seconds. The actual number of slots can range between 60 and 15360 and must be divisible by 60. The number of slots (M1) is typically set at 4500.

Transmission synchronisation is provided by UTC time and the start of a group of M1/60 slots always begins at the boundary between two successive UTC seconds.

In turn, this requires access to a UTC time source of which GPS is a typical source. In fallback modes, a VDL4 station can reference timing information provided by another station or by monitoring the transmissions of others. However, the primary source of timing information, and the only one specified for use other than in failure modes, is a UTC time source.

A transmission is either wholly within a single slot or occupies a series of contiguous slots. Each transmission comprises the necessary power-up and power-down transmissions and the “burst format” derived from ISO 3309 (HDLC). Note that in VDL4 terminology each frame is a burst and so there may be several bursts in a transmission.

3.2.2 Burst Format

The VDL4 Burst Format is shown in Figure 1 (reproduced from the VDL4 Technical Manual). The framing is standard HDLC, with HDLC flag characters and a CRC based Frame Check Sequence (FCS). The FCS is (as with VDL2) the standard 16-bit CRC specified in ISO 3309. Unlike VDL Mode 2 there is no Forward Error Correction.

Description	**Octet**	**Bit number**							
		8	7	6	5	4	3	2	1
flag	-	0	1	1	1	1	1	1	0
TCP change flag, reservation ID (rid), version number (ver)	1	s_{27}	s_{26}	s_{25}	ver_3	ver_2	ver_1	rid	tc
source address (s)	2	s_{24}	s_{23}	s_{22}	s_{21}	s_{20}	s_{19}	s_{18}	s_{17}
	3	s_{16}	s_{15}	s_{14}	s_{13}	s_{12}	s_{11}	s_{10}	s_9
	4	s_8	s_7	s_6	s_5	s_4	s_3	s_2	s_1
message ID (mi)	5	in_k	mi_k	. . . . . .		mi_4	mi_3	mi_2	mi_1
information	6								
	7 - n-5			. . . .					
	n-4								
reservation data (rd)	n-3		in_1	rd_k	. . . . . .				
extended reservation ID (erid)	n-2	$erid_k$	. . . . .			$erid_1$			rd_1
CRC (c)	n-1	c_9	c_{10}	c_{11}	c_{12}	c_{13}	c_{14}	c_{15}	c_{16}
	n	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8
flag	-	0	1	1	1	1	1	1	0

Figure 1 Burst Format

The standard frame format comprises the Source Address, a Message Type Identifier, an Information (payload) field, and fields to support slot reservation. It should be noted that the “TCP flag” has got nothing to do with the IETF’s Transmission Control Protocol and is used to indicate a “directed synchronisation burst” i.e. ADS-B information sent to a specific destination. The destination address (if not a broadcast) is encoded into the reservation request (typically a reservation for a response to the transmission, but a null reservation is also permitted).

The different Message Types are mostly concerned with the slot reservation schemes, although they can also indicate different types of information field as well.

3.2.3 Slot Management

The VDL4 specification requires each to station to keep slot reservation information for more than the next four SuperFrames. This information indicates whether a given slot has no reservation associated with it or whether it is reserved for broadcast or point-to-point communications, who by and, in the latter case, to where. This information is built up through monitoring the channel on net entry and maintained as long as a station is part of a VDL4 network.

A slot that is not reserved is available to any station that wants to use it (provided a transmission in the previous slot has ceased), and a station needing to transmit may do so using the *Random Access Procedures*, as soon as the slot comes round. This procedure uses a *p-persistent* algorithm to determine the probability of transmission, plus listen before send (i.e. a form of CSMA).

Of course, if two or more stations transmit in the same slot they will probably garble each other's transmissions, which limits the utility of the Random Access approach.

In order to improve upon the performance of Random Access, VDL4 provides slot reservation schemes, of which the underlying principle is that when a station transmits, it may include a "reservation" for a later slot or a series of contiguous slots. Such a reservation may be either for itself or another station, and may even be on a different frequency.

Other stations will see that reservation and include it in their own Reservation Tables. They should then avoid transmitting in such a reserved slot and thus avoid the performance loss that would otherwise occur as a result of "slot collisions".

3.2.4 Slot Reservation

Slot reservations are either made for specific transmissions (usually point-to-point) or for "Periodic Broadcasts".

In the former case, a transmission, which may well include user data, includes a reservation for a future transmission. A station may transmit a reservation in a single unused slot reserving a block of slots later on in order to send a large message while minimising the risk of a collision. When sending that message, it may also reserve a slot for the recipient of the message to respond in. That way, an efficient message exchange can take place.

In the latter case, the reservation is for a repeated slot used for a broadcast transmission at a regular interval. This transmission is expected to be used for synchronisation (ADS-B) data and may also be used to reserve other slots for data transfer.

Additionally, on channels reserved as Global Signalling Channels, a set number of slots after each UTC second may be reserved for the exclusive use of Ground Stations.

When multiple frequencies are in use, it is also possible to "autotune" a response i.e. to specify a different channel on which the response reservation has been made.

3.2.5 Network Entry

Before a VDL4 station can start transmitting it has to somehow "join the net". One of the simplest ways to do this is to monitor other station's transmissions for one SuperFrame (i.e. a minute) and then transmit in an unreserved slot. Such a transmission will typically include a Synchronisation burst reporting the station's 4D position and will include a Periodic Reservation thus establishing the station in the "net".

This procedure is satisfactory provided that there is no need to transmit on initial net entry and there are sufficient unreserved slots available. To enter a net earlier (than one SuperFrame), there are two mechanisms available. Either the station can send a "plea request" to a station that it has identified as already being in the network, or it can use a "Big Negative Dither" reservation request to reserve slots far enough away in the SuperFrame to have not already been reserved.

In both cases, the reservation request is sent in a slot for which the sender does not have positive information that it is reserved.

A plea is hopefully responded to with the offer of an already reserved slot in a relatively short time, while a BND forcefully gets a slot at a later time. Either way, the station should have gain access to a number of slots reserved for itself.

The initial transmission uses a variation of Random Access known as *Delayed Random Access*. In this mode of operation, the “Listen Before Send” procedure is delayed 4ms after the start of the slot. This minimises the risk of a station with only a partial slot map interfering with an existing reservation.

3.2.6 Fair Allocation of Bandwidth

VDL4 provides limits on how often a station can reserve slots for Periodic Broadcast. It also requires Periodic Broadcast slots to be regularly changed. Such mechanisms have the effect that one station cannot exclude others for ADS-B purposes. However, reservations for data transmission slots are not limited by the protocol or procedures, and an “aggressive” slot grabbing station could gain unfair access to the channel.

There is also no priority based mechanism for pre-empting channel reservations. Once a reservation has been made, another station cannot transmit a pre-empting reservation at a higher priority. While VDL4 does specify priority based queue management in each station and for peer-to-peer communications, unlike VDL Mode 3, station transmissions are not prioritised with respect to each other. It is thus similar to VDL Mode 2 in this respect.

The case has been made before for assigning different priority data to different channels. However, this assumes that channels are readily available and could, for example, mean that a high priority channel is congested while a low priority one is not.

3.2.7 Slot Stealing

If there are insufficient unreserved slots, it is possible for a VDL4 station to “steal” another station’s slot provided that it is far enough away. In turn this depends upon receipt of ADS-B information so that each station knows where the others are.

The simplest and most robust case of slot stealing occurs when a slot has been reserved by “Station B” to send a message to Station “D”. “Station A” can use this slot to send a message to “Station C” provided B and D are sufficiently far away from A and C for the transmissions not to interfere with each other. The SARPs define a distance metric designed to ensure that this is the case.

It is possible to miss a reservation made by “Station B” during such a transmission and to prevent other stations from receiving such a reservation. This can lead to later collisions but as such reservations will typically be in the context of Station B to D communications, this should not be an issue.

It is also possible to steal a broadcast slot. The victim of this form of crime has to be further away than in the former case but that is the only constraint. The downside of such a “steal” is that it stops other (nearby) stations from receiving the victim’s ADS-B information and any slot reservations made in such a transmission. This makes future collisions more likely. A particular problem will result from a station intermediate between the stealing station and the victim, which will then not be aware of any reservations made by the victim and may then view such slots as unreserved – only to transmit in them resulting in a garbled transmission.

3.2.8 The Impact of Lost Reservations

VDL Mode 4’s reservation scheme can work well and add significantly to performance. However, there is no acknowledgement protocol associated with reservations and every reservation request is effectively broadcast to all stations. It may be garbled by another transmission or rejected by a receiving station due to an error induced by background noise.

It may be that only a subset of all stations in a network receive without error a given reservation request – another downside of the lack of any Forward Error Correction.

This means that a slot that has been apparently reserved is viewed as unreserved by other stations in the network. As the *Random Access* procedures require “listen before send”, this should not result in collisions due to a Random Access transmission. However, it may result in another station reserving the same slot.

Contention resolution procedures are defined in the draft SARPs to handle this situation when recognised by competing stations. The conflict resolution procedures avoid some conflicts but do not appear to resolve all conflicts. The procedures do also take transmission priority into account even though the specification does not appear to allow for deliberate slot pre-emption.

Of course, the conflict resolution procedures are not effective when one or both stations are unaware of their competing reservations.

3.3 VDL Mode 4 Data Link Layer

The VDL4 data link layer is derived from the AVLC specified for VDL Mode 2. In turn, this is a version of HDLC.

For air-air communications, a connectionless data transfer service is supported using HDLC Unnumbered Information (UI) frames.

For Air/Ground Communications, connection mode HDLC data transfer principles are utilised to provide a reliable data transfer service optimised for operation with VDL Mode 4.

3.3.1 Air/Ground Data Communications

The Link Establishment (LE) and Handoff (HO) procedures follow the VDL Mode 2 approach of using an exchange of XID frames to establish a datalink and to perform handoff to another Ground Station. Data Transfer uses I-Frames with RR and SREJ frames to provide flow control and reliable data transfer. Special procedures are used for the transfer of both data and XIDs in order to take full advantage of VDL4 slot reservation protocols.

Air/Ground communication is connection mode and, as with VDL Mode 2, an air-initiated XID exchange is required before data transfer can begin.

3.3.2 XID Exchange

XID exchanges are used as in VDL Mode 2 to advertise Ground Station capabilities (i.e. GSIFs), to establish a Data Link and for Handoff from one Ground Station to another. However, XID messages can be long transmissions and to support this whilst taking maximum advantage of the reservations protocols, the RTX burst is specified.

This is a short single slot message sent to the recipient to request the sending of an XID message. The RTX burst also contains a slot reservation for the response. The response to an XID will reserve slots for the XID which can then be sent in those reserved slots.

The RTX is repeated if there is no response from the Ground Station.

Additionally, a Ground Station can use a periodic broadcast and its reserved slots to uplink XIDs such as a GSIF.

3.3.3 Data Transfer

3.3.3.1 Short Message Transfer

Short messages (typically up to 83 octets or three slots) are transferred as a single one-off burst. This may be in a slot reserved by (e.g.) a periodic broadcast (using incremental reservation procedures) or transmitted in an unreserved slot. The burst contains the I-Frame containing the information to be transferred and also contains a reservation for a response from the destination station.

The response should contain an RR or SREF frame, or a protocol error response (e.g. DM or FRMR). If no response is received in the reserved slot then the original transmission is repeated as described above. After ‘n’ retries with no response, the transmission attempt fails.

Note that retransmission on a MAC layer timer appears to be used to perform recovery on packet loss rather than an AVLC timer. However, the specification is not clear on this issue.

3.3.3.2 Long Message Transfer

Longer messages start with an “RTS” burst. This is a short single slot message sent to the recipient to request the sending of a long message. It includes information on the priority and number of frames in the message. The RTS burst also contains a slot reservation for the response.

The response will typically include an RR and a reservation sufficiently large for the RTS sender to send at least the first frame. Once this has been returned, the RTS sender can now send as many I-frames as the reserved group of slots permit. In turn, this transmission also includes a reservation for a response. The response will typically include an RR (to acknowledge error free receipt) or an SREJ (to indicate an error), and may also contain a reservation for the transmission of further I-frames as indicated in the RTS.

This procedure includes retransmission of the RTS in the event of no response and takes full advantage of the VDL4 reservation protocols in order to optimise performance.

Note that the response to a retransmitted RTS can include an RR that acknowledges receipt of the previous transmission (indicating that it was the response that was lost). In which case the reserved slots are either unused to used to transmit further data frames.

As with Short Message Transfer, retransmission on a MAC layer timer appears to be used to perform recovery on packet loss rather than an AVLC timer. However, the specification is not clear on this issue.

3.4 Data Link Operation

In summary, and for Air/Ground communications, VDL4 operates as follows:

1. An aircraft performs the Net Entry procedures after monitoring a channel for a given period of time. During that time it may also receive one or more GSIFs from Ground Stations as broadcast XIDs.

2. Typically using position information, it chooses the “best” Ground Station and sends it an RTX message in an unreserved slot. The RTX includes a reservation for the Ground Station’s response.

3. The Ground Station responds to the RTX and reserves a suitable slot for the aircraft’s response. This may be on the same or a different frequency depending on current channel utilisation.

4. The aircraft sends its XID_CMD_LE to initiate a data link, and reserves a response slot for the Ground Station's XID_RSP_LE.

5. The Ground Station returns the XID_RSP_LE and data transfer can commence.

6. The data transfer procedures may then take place as described above.

Later on, the aircraft may find it difficult to keep in contact with the Ground Station as the signal strength drops off and interference increases. Either due to such effects or simply because a significantly nearer Ground Station has been detected (its GSIF has been received), the Handoff procedures are invoked.

As described above, the aircraft sends an RTX to the new Ground Station in order to initiate the exchange of the XID_CMD_HO and the XID_RSP_HO. This exchange performs the Handoff procedure as specified by VDL Mode 2 and the aircraft's datalink is now routed through the new Ground Station.

3.5 Conclusion

This analysis has identified a number of advantages of the VDL Mode 4 specification compared with VDL Mode 2 and a number of disadvantages. The big issue though is the performance of the VDL Mode 4 network as a data communications network rather than as a surveillance network.

3.5.1 Advantages

- Slot Stealing of point-to-point transmissions provides a useful mechanism for avoiding the performance problems due to "Exposed Transmitters". This is a known problem in ACARS and VDL Mode 2, and should give a performance benefit. It is also taking full advantage of the better noise immunity of GFSK when compared to VDL Mode 2's D8PSK modulation scheme.

- The same mechanism should also permit frequency re-use by Ground Stations over shorter distances than with VDL Mode 2, but this will demand better co-ordination of Ground Station transmissions in order to avoid Hidden Terminal effects.

3.5.2 Disadvantages

- The lack of Forward Error Correction limits the performance of data transfer and increases the risk of conflicting reservations.

- Slot stealing of broadcast transmissions has a quiet different effect to slot stealing of point to point transmissions. When applied, it decreases the useful size of the "communications cell" and effectively provides an upper bound on cell diameter and the numbers of aircraft in each cell. While it can also be viewed as mitigating the Exposed Transmitter problem there are additional effects resulting from missed reservations that are undesirable consequences of the procedure.

- The lack of priority based pre-emption of slot reservations is also an issue for CNS/ATM Applications. While each station's transmission queue is prioritised, MAC Layer contention for access to the channel by different stations is not prioritised. This is not dissimilar to the situation with VDL Mode 2 where implementations can locally prioritise the transmission queue, while there is no pre-emptive access to the channel as a direct consequence of CSMA.

3.5.3 Performance

Under VDL Mode 4, it will be expected that the access delay increases with the number of transmitting stations, due to the proportion of reserved slots increasing and a longer delay being experienced before slots become available. Given that some of the transmissions are performed under Random Access procedures (using apparently unreserved slots) the number of collisions will also rise as the number of transmitting stations increases resulting in a need for retransmissions thus further increasing the load on the channel. The increase in access delay would thus be expected to be exponential rather than linear.

As the load further increases the point is reached where the cell size shrinks and the time to transmit to more distant stations increases even more greatly.

Measurement of transmission power levels is also very important here as each station has to decide when a slot is empty. A certain level of background noise has to be permitted but if the permitted noise level is set too high then remote transmissions (and hence slot reservations) will not be received, simultaneous transmissions will occur, resulting in garbling for intermediate stations.

The net result is a non-deterministic transit delay which grows exponentially with load – much the same characteristic as is exhibited by the CSMA scheme used by VDL Mode 2. As with VDL Mode 2, VDL Mode 4 as currently defined thus has limited suitability for time critical CPDLC applications and the number of aircraft supported by a given Ground Station will have to be artificially limited (as with VDL Mode 2). This may be performed by the Ground Station rejecting Link Establishment or Handoff requests, or autotuning an aircraft to an alternative frequency, once a set limit has been reached. As such limits will need to be conservatively set to meet ATC demands, there will be an overall under-utilisation of capacity.

However, the RTS procedure does seem to offer a way of improving upon this situation. When an aircraft sends an RTS to a Ground Station, the Ground Station gets information on message priority and can schedule the aircraft's data transmission taking into account requests from other aircraft. Thus a Ground Station can introduce a more deterministic performance characteristic by applying such a procedure.

The RTS procedure is specified to be symmetric, thus each aircraft has to be relied upon to correctly schedule Ground Station uplinks. This does not seem so appropriate as the Ground Station is perfectly capable of doing this itself, using its protected slots to reserve slots for its own uplinks, scheduling according to data priority. In contrast, a poorly performing aircraft implementation could lower the overall performance of the network.

Airborne implementations are best kept simple given the certification requirements and long equipment lifetimes. In contrast, Ground Systems are more flexible and can rapidly evolve as experience in the system's operation grows. Ground control would thus seem the most appropriate strategy.

It is believed that a more ground oriented approach to slot management would:

- Allow a greater proportion of slots to be used for data transfer.

- Allow priority based allocation of slots to be introduced thus permitting greater capacity utilisation whilst maintaining performance guarantees for CNS/ATM Applications.

- Permit a more rapid introduction of better slot management algorithms.

4. The Frame Mode SNDCF

4.1 Background

The first generation of ICAO air/ground datalinks all used ITU-T recommendation X.25 as their network access protocol. This provided a common service interface using an industry standard, but is not necessarily the most efficient way to provide access to a mobile subnetwork. This is particularly true of the ICAO VHF Digital Link (VDL). VHF communications are largely constrained to line of sight operations and to maintain communication while in flight, an aircraft must necessarily switch (i.e. Handoff) from one VDL Ground Station to another. Maintaining X.25 virtual circuits across a Handoff has proved to be a particular complex procedure with significant overhead.

Specifically, with X.25:

1. The ISO 8208 packet layer provides for multiplexing and flow control of multiple data streams (virtual circuits), each of which may have a different priority.

2. To have transfers at different priorities a separate virtual circuit has to be established for each priority.

3. A virtual circuit adds overhead in terms of both packet headers and control messages.

4. Under VDL Mode 2, existing virtual circuits must be cleared and new virtual circuits established every time an aircraft changes ground station.

5. The ISO 8208 packet layer is an additional and complex software component.

In support of the first generation of ICAO air/ground datalinks, the ATN SARPs define the Mobile X.25 SNDCF. This is an adaptation of the existing X.25 SNDCF specified in ISO/IEC 8473-2 to support the exchange of CLNP packets over X.25 subnetwork connections. The adaptation supports the negotiation and use of the ATN specified data compression algorithms for packets transferred over such an X.25 subnetwork connection.

In support of second generation Air/Ground datalinks, the ATN SARPs specifies a more generic SNDCF that has become known as the Frame Mode SNDCF (this name arose from the original purpose of supporting VDL Mode 3 Frame Mode communications). This is designed for lightweight datalinks and may be used directly over a datalink or MAC layer service. Its functionality has also included lessons learnt from the deployment of the Mobile X.25 SNDCF and, in particular, is more extensible, exploits more of the capabilities of Deflate, and allows for new features to be introduced in a backwards compatible manner. All new datalinks adopted by the ATN are expected to use the Frame Mode SNDCF.

4.2 Characteristics

The “Frame Mode SNDCF” comprises the following functions:

- An Air/Ground Communications Sublayer (A/GCS) to provide for multiplexing of different data streams and support of stream mode data compression algorithms (e.g. Deflate).

- Deflate based Stream Mode Data compression

- LREF based CLNP compression

- CLNP Header reformatting (to improve Deflate compression).

The resulting SNDCF has the following characteristics:

1. It can operate over any data communications service providing a low probability of packet loss, duplication, corruption or re-ordering, including the VDL Mode 3 Frame Mode Service, the VDL Mode 2 AVLC, and ground-ground services such as Frame Relay.

2. It incorporates features known to be missing from the current Mobile SNDCF concerned with the introduction of new functions in a backwards compatible manner.

3. It supports maintenance of the Data Compression State across data link Handoffs.

4. It includes optional security features for data link authentication.

5. Deflate Dictionaries and re-synchronisation are supported.

6. Other Routable protocols such as IPv4 can also be supported concurrent with CLNP.

4.3 Design Approach to a new SNDCF

The design of the original Mobile SNDCF was predicated on an ISO 8208 base and is heavily influenced by it. The Frame Mode SNDCF is predicated on a much more basic service, that of a simple packet mode service that can be either:

- a connection mode datalink,
- a singlecast connectionless service, or
- a broadcast connectionless service.

The specification for both Deflate and LREF compression permits recovery from data loss. This permits operation over a connectionless service and performance will be satisfactory provided the probability of data loss or re-ordering is low. The same mechanism also supports operation over a reliable connection mode service and recovery from a connection reset.

Deflate is not suitable for use over broadcast datalinks. However, an adaption of LREF is supported for this purpose.

However, even with a connectionless service there will still be a need to negotiate compression options and to provide for recovery in the event of data corruption. In short, there always needs to be some aspects of a connection mode service in order to build and maintain a relationship between compressor and decompressor. This is especially true of Deflate, although it is possible to operate LREF in a true connectionless fashion (i.e. for broadcast communications) – this is only possible by regularly sending uncompressed CLNP PDUs.

When a connectionless service only is available and the service is known to have an unacceptably high rate of packet loss, then some sort of frame level acknowledgement will be required with retransmission in the event of no acknowledgement being received.

The specification of Deflate also permits:

- Use of Deflate Dictionaries in order to speed the convergence of the compression algorithm to the optimal compression ratio.
- Recovery from checksum errors by returning to the last received packet rather than restarting the compressor.

- Re-use of compression state after Ground Station Handoff (where this is a feature of the datalink).

4.4 Datalink Management

The Frame Mode SNDCF provides a Data Link Management protocol in order to negotiate compression options and to maintain the data communications path.

Modern communications networks tend to use an "out-of-band" approach to data link management. This is true of Frame Relay, Asynchronous Transfer Mode and ISDN. This technique ensures independence between data link management and data transfer and enables a very lightweight protocol approach to data transfer. A similar approach to the data link management protocol is specified here. This simplifies the specification and is bandwidth efficient.

In order to provide for a Datalink Management protocol while maintaining lightweight communications, the Frame Mode SNDCF adopts the following strategy:

1. The air/ground data stream is subdivided into a number of bi-directional "logical channels", each identified by a channel number.

2. Data packets including data link management packets are always sent in the context of a logical channel i.e. they are encapsulated with a simple header identifying the channel and priority.

3. Subject to maintaining the semantics of data priority, packets sent on different channels may be concatenated and sent as a single transmission frame.

4. A Data Link Control Protocol (DLCP) is specified in order to negotiate data link capabilities and compression options, and to manage the purpose and use of each channel. Channel zero is reserved for the DLCP.

5. A channel assignment can specify that the data packets are Deflate compressed before transmission on a given channel. It is also possible to have multiple Deflate compressors with different groups of channels assigned to different compressors.

The channel concept firstly allows the separation out of the DLCP from user data and secondly allows for different Deflate compressed streams to be multiplexed together. It is also extensible as it permits other compression protocols to be introduced later (e.g. ADCMP for audio compression) and used in parallel to Deflate on other channels. Some channels (in addition to channel zero) could also be uncompressed, if that was needed.

The specification also deliberately groups multiple channels together for compression as a single data stream. This is done to ensure the most efficient compression while still allowing the channel concept to be used to differentiate different data streams. Essentially, a Deflate Compressor becomes a Server asked to compress a packet before it is encapsulated with a channel header and appended to the transmission frame. The channel determines the choice of compression Server.

At least two user data formats are foreseen. These are ISH PDUs, and CLNP PDUs that cannot be LREF compressed – or more generally any ISO TR 9577 identifiable protocol – and LREF compressed CLNP PDUs. However, the channel concept can readily extend to supporting other data formats including IP and even ACARS messages.

A channel is not associated with priority, which is totally independent of the channel concept. Essentially a channel is a dynamically assigned label that identifies a type of data (e.g. CLNP) and links it to a compression algorithm.

4.5 Architecture

The architecture of the proposed approach is really more than an SNDCF. As illustrated in Figure 2, the specification is in two parts: the SNDCF proper and an Air/Ground Communications Sublayer (A/GCS). The A/GCS comprises the multiplexing of many logical channels, the DLCP and the Deflate compression (which operates across channels). The SNDCF itself is responsible for providing the SN-Service over the A/GCS and also incorporates LREF compression, which is CLNP specific.

There may be other users of the A/GCS in addition to the Frame Mode SNDCF. For example, if IPv4 packets are also conveyed then IPv4 will be a separate user of the A/GCS.

The underlying Data Link Layer is assumed to provide:

- A communications service that provides sufficiently[1] reliable delivery and preservation of frame transfer sequence order.
- an optional broadcast ground to air unacknowledged connectionless service.
- “Join” and “Leave” events to signal contact with a new Ground Station/Aircraft, and loss of contact, respectively.

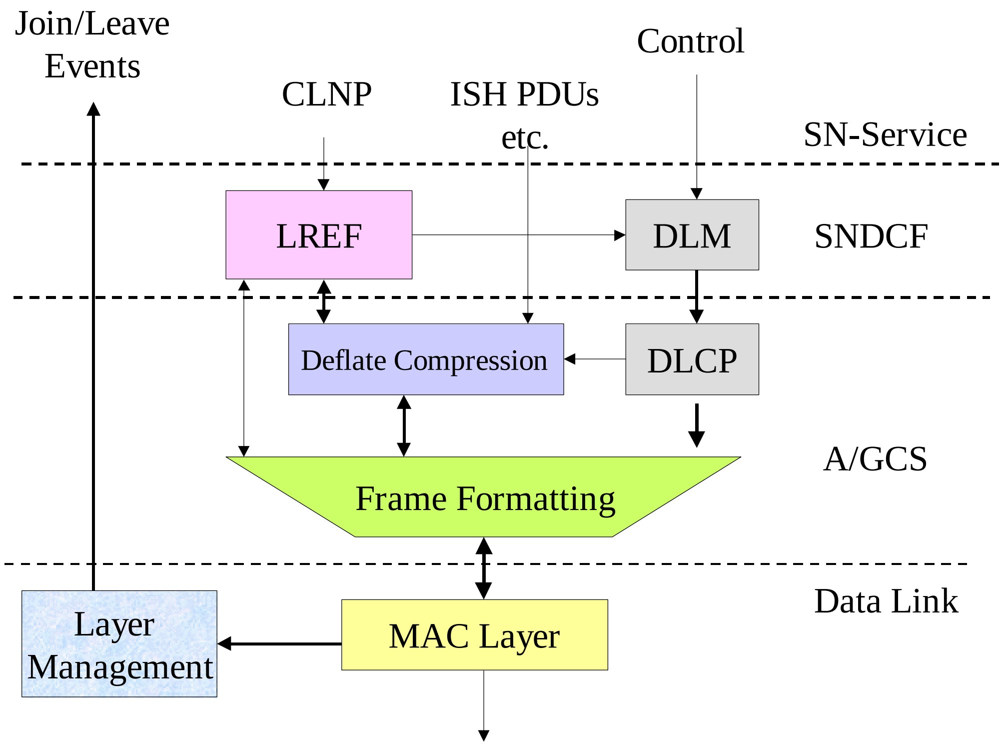

Figure 2 Frame Mode SNDCF Architecture

The A/GCS provides a (DLCP supported) Data Link Management service to its user – the SNDCF’s Data Link Manager (DLM) - and a data communications service. The Data Communications service is channel oriented, with channel assignment performed through the Data Link Management interface. The data communications service is packet oriented

[1] Sufficient to provide to meet the overall throughput requirements.

and is guaranteed to maintain packet delivery order but does not guarantee to deliver all packets. The probability of packet loss is small but can come about (e.g.) through detection/recovery from a Deflate checksum failure.

The SNDCF provides the required SN-Service. It will use the A/GCS Data Link Management Service to maintain at least one channel for ISH PDUs, and CLNP PDUs that cannot be compressed using LREF. Otherwise, the SNDCF will use the A/GCS to assign a new channel for each local reference needed and use this channel to transfer LREF compressed CLNP PDUs associated with that Local Reference.

In ATN architecture, the IS-SME receives the Join/Leave Events and issues commands to the DLM.

4.6 Understanding the Frame Mode SNDCF

4.6.1 Channels

A channel is simply a labelled flow of data. The label is a 12 bit number (the channel number) and, initially, zero is the only defined label. Channel 0 is reserved for the DLCP, and the DLCP is used to assign further channel.

Each channel constitutes an agreement to transfer a certain data type and to apply a specific compression algorithm to that data (or to send it uncompressed). There are no additional semantics associated with a channel, nor is there any flow control associated with a channel.

A Frame may consist of more than one data packet, each on a different channel, concatenated together. It is thus possible for a DLCP message defining the use of a channel to be immediately followed by the first data packet that uses that channel number. Indeed, this is how it is meant to be used with a channel only defined when it needs to be used.

Each such data packet may also have a priority associated with it. There is no requirement for all data packets on the same channel to have the same priority (because there is no flow control or data recovery done at the channel level), but there are rules on when packets with different priorities can be concatenated together into the same transmission frame. These rules are there to ensure that higher priority data is not denied access to the medium by lower priority data.

When the datalink imposes a maximum frame size, it can be difficult to calculate the maximum SNSDU size as this depends on the achieved compression ratio. It is thus recommended that the maximum SNSDU size is calculated assuming no compression and that concatenation is relied upon to take full advantage of compression by transferring more data per frame.

4.6.2 Compression

Deflate compression is an integral part of the A/GCS specification. However, it is not the only data compression algorithm that can be supported. The specification is deliberately designed as extensible to enable other data compression algorithms to be introduced later on in a backwards compatible fashion.

The basic concept is that a sender has one or more Deflate compressors associated with each Peer System and that each compressor may be associated with one or more channels. The reason for allowing for more than one compressor per Peer System is that each may be initialised with a different “dictionary” that allows compression to rapidly converge on an optimal compression state for a specific type of data.

In this context, a “dictionary” is simply a well known string (less than 32KB long) of frequently occurring symbols. It is fed into the compressor (and the output discarded) when the compressor is initialised; the same string is similarly feed into the decompressor in the Peer System. Provided that the symbols in the dictionary re-occur less than 32KB into the data stream (from where they occur in the dictionary) they will be immediately compressed on transmission.

When a channel is assigned, a Deflate compressor may be associated with it. The actual compressor is identified by a Dictionary Identifier that is known *a p*riori by both sender and receiver.

As Deflate is a stream compression algorithm it is very important that the decompressor works on an identical data stream to the compressor. Compressed packets are thus decompressed in strict transmission order and the underlying datalink is assumed not to re-order packets. If re-ordering does take place then a decompression error will result.

In order to provide for error detection and recovery, each compressed packet is appended with a checksum taken on the uncompressed data. If this checksum cannot be verified on reception, the data is discarded and a Link Reset message returned (using the DLCP) identifying the offset into the data stream of the last correctly received byte. Compressed data on the affected streams continues to be discarded until the compressor acknowledges the reset and restarts from the identified stream offset.

This procedure allows for recovery from the occasional re-ordering or packet loss. However, throughput will suffer if this event is more than infrequent.

Note that only the leading window edge of the compressor and decompressor have to be synchronised. The reset procedure typically results in the decompressor having a larger history window than the compressor. This is not a problem because the compressor cannot make a backwards reference beyond its own history window and hence all backwards references will be within the decompressor's history window. A problem will only result if the decompressor has a smaller history window that the compressor - which can lead to it receiving a backwards reference that it cannot resolve. The link reset procedures are designed to avoid this situation occurring.

4.6.3 Security

The A/GCS optionally provides a framework for authentication of the source of each frame.

The DLCP provides for the exchange of security related information on link establishment which permits a shared encryption key to be agreed. This may then be used to compute an encrypted checksum for each transmission frame, which may be appended to the concatenated data packets in each frame. A receiver can use this to verify the source and integrity of the data.

4.6.4 The DLCP

The DLCP is a channel management protocol and does not itself transfer any user data. Its functions are:

- Data Link Initialisation
- The assignment of channel numbers and the semantics of the data carried over a channel.
- Compression algorithm Management, and
- The optional authentication and verification procedures.
- The backwards compatible introduction of new features and compression algorithms.

The protocol is designed to be robust (i.e. it can recover from data loss, duplication, etc.) but is also lightweight and avoids acknowledgements unless necessary. For example, Channel Assignment is unacknowledged; however, if this message is lost then the peer system will respond with an error when data is sent on an unassigned channel.

On data link initialisation, information on the capabilities of each peer system are exchanged including the compression algorithms supported, with each Deflate Dictionary supported being treated as a separately identified compression algorithm.

A channel is simply allocated by a "Channel Start" message identifying the channel number, the data type associated with the channel and an optional compression identifier. Data types are currently defined for:

- ISO TR 9577 formatted NPDUs (e.g. ES-IS and uncompressed CLNP)
- LREF compressed CLNP
- CLNP NPDUs with reformatted headers (for optimal Deflate compression)
- IPv4 packets.

Channels are symmetric once assigned i.e. have the same data type and compression algorithm associated with them in each direction. To avoid conflicts when allocating channel numbers, each side starts with either high or low channel numbers and at least one channel must remain unassigned.

There is no requirement for a data type to be associated exclusively with a single channel. For example, LREF compressed data may be associated with different channels each themselves associated with a different compression algorithm (or Deflate Dictionary). The sender can then send data with different Traffic Types using the most appropriate compressor for that Traffic Type. As local references are associated with the datalink rather than the channel, there is no loss in LREF efficiency here.

Channel numbers may be de-assigned with a Channel Deallocate message and reset by a Channel Reset message. The purpose of the latter message is really data type specific and may be used, for example, by LREF error recovery procedures.

The DLCP also supports retention of compression state and channel assignments across Ground Station Handoffs. Each datalink session is labelled on initialisation by a Ground Station and, on Handoff the old session identifier is presented to the new Ground Station which will recover the compression state and channel assignments if possible.

This technique avoids an aircraft having to keep track of whether or not a Handoff is to a new GNI or not. In both cases, it simply presents the old session identifier and it's up to the Ground User to recognise whether the compression state information is local or whether it has to be retrieved from another Air/Ground Router.

It should be noted that the DLCP specifically requires that an aircraft initiates the data link procedures. This is to optimise the procedures for retention of compression state on handoff.

Extensibility is managed by requiring a receiver of the datalink initiation message to ignore unrecognised parameters. This permits new parameters corresponding to new facilities to be introduced without affecting existing implementations. When a new parameter in the datalink initiation is ignored by the receiver (i.e. the response does not include a response to such a parameter), then the initiator assumes that it is communicating with an older implementation and adjusts its procedures accordingly.

5. Accessing VDL Mode 4 Using the Frame Mode SNDCF

5.1 Data Link Establishment

A “Join” Event from the datalink is the trigger that causes the SNDCF Datalink initiation sequence to be invoked. This event takes place on an aircraft and the “Join” Event is sent to the “IS-SME” (the function specified in the ATN SARPs for mobile datalink management). The IS-SME then instructs the SNDCF DLM to send a DLCP Data Link Start (DLS) packet to the Ground Station side using the datalink.

The “Join” Event can occur after the datalink has been established by the data link layer service, or it can indicate the possibility of a datalink which is only initiated when an attempt is made to send the first packet over the datalink. When it occurs is a datalink issue rather than an issue for the SNDCF.

If and when the Frame Mode SNDCF is used over the VDL Mode 2 AVLC, it is expected that the Join Event will occur before the datalink is established i.e. when the GSIF is received from a suitable Ground Station. This gives the opportunity to include the DLS as an XID parameter in the XID_CMD_LE. Similarly, the response to the DLS can be included as a parameter to the XID_RSP_LE. This is a bandwidth efficient means of establishing the datalink in VDL Mode 2.

It would also be possible, under VDL Mode 2 to have delayed the Join Event until the AVLC data link was established. However, this requires the exchange of the DLS frames as AVLC I-Frames thus doubling the number of Air/Ground Transactions to initiate a datalink. As the performance of VDL Mode 2 is more dependent on the number of messages sent rather than their size, it is clearly a much better strategy to send the DLS as an XID parameter.

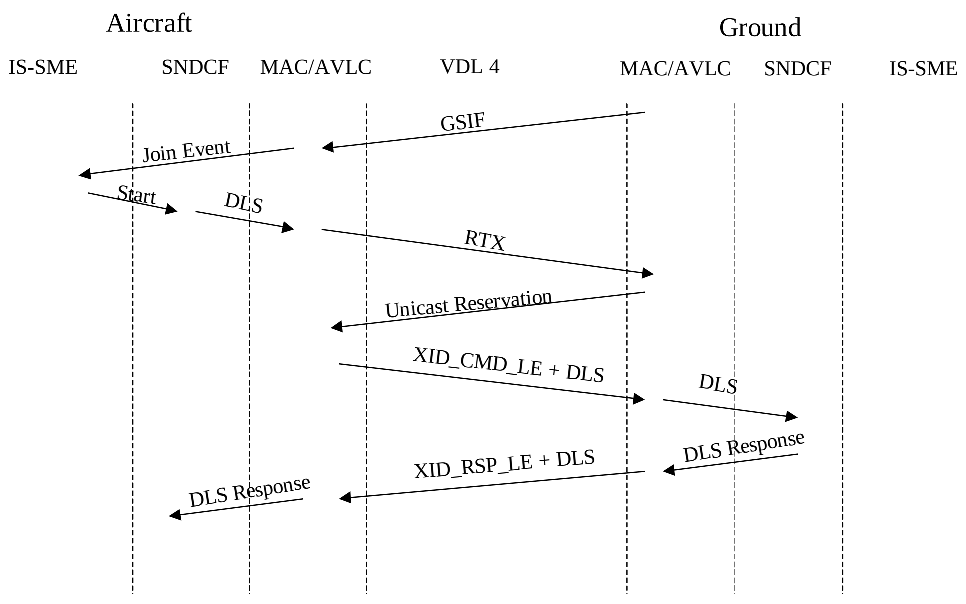

Figure 3 Data Link Initiations - Sequence of Events

Under VDL Mode 4, the same approach can be adopted. However, there is an extra issue here – the lack of Forward Error Correction (FEC). Without FEC, there is a much greater probability of data corruption which increases with frame size. Sending a large XID frame may thus lead to an increased probability of retransmission thus reducing the benefits of combining XID and DLS. Also VDL4 performance is believed to be less dependent on the number of transmissions than VDL Mode 2.

Hence, it may be more appropriate with VDL4 to delay the Join Event until after the datalink has been established. This is an issue that can only be completely resolved by validation and flight trails and both approaches should be catered for in a specification.

Figure 3 illustrates the sequence of events needed to establish a data link, assuming that the XID and Data Link Start (DLS) packets are merged. Clearly, the sequence becomes extended with an RTS exchange if the Join event comes after establishment of the datalink.

5.2 Data Transfer Options

The Frame Mode SNDCF A/GCS queues frames for transmission and attempts to concatenate as many frames as possible into the same transmission. It requires a datalink service that is generally reliable (i.e. will deliver frames without error or duplication and in sequence), but can recover from errors, albeit with a loss of performance.

5.2.1 The Connectionless Data Communications Service

In VDL4, the connectionless data transfer service offered by the Unnumbered Information (UI) frame format could be used to transfer A/GCS frames, especially as UI frames are acknowledged by Unnumbered Acknowledgements (UA).

However, when a response to a UI frame is not received there is no way of knowing whether it was the response that was lost or the original frame. Retransmission on lack of response can thus result in a duplicate frame being received and processed as there is no way that the recipient can detect this as a retransmission. Processing the duplicate frame will cause a decompressor error and the consequential overhead of resynchronisation.

5.2.2 The VDL Mode 4 AVLC

Alternatively, there is the VDL4 AVLC implementation which is anyway specified for Air/Ground use. This permits a station to send one or more I-frames concatenated into a single transmission. The response to this is either an RR acknowledging receipt of all transmitted I-Frames or an SREJ identifying any received in error. With no FEC, it makes sense to concatenate individual small frames into a single transmission rather than the data in those frames as this minimises the retransmission overhead.

As each frame is numbered by the AVLC, retransmissions can be identified by the recipient and discarded if the original had already been correctly received. This mode of transmission thus avoids the message duplication problem associated with UI frames. It also efficiently permits several frames to be sent in one transmission without increasing the risk of total loss due to a single error.

This advantage will be lost if the A/GCS concatenates many data packets itself into a single transmission frame, and hence this behaviour will need to be modified for VDL4. Otherwise, the VDL4 data transfer mechanism appears to be well suited to ATN data transfer.

Figure 4 illustrates the sequence of events for data transfer using the Long Message Transfer protocol and assuming no communications errors.

5.2.3 A Simplified Data Link Protocol

An intermediate data transfer mechanism could be proposed (i.e. intermediate between a UI/UA exchange and the full AVLC) by defining a simple numbered acknowledgement scheme i.e. each transmission is numbered and so is each acknowledgement. This would avoid the packet duplication problem inherent in the UI/UA exchange. Such a scheme would also take full advantage of the existing MAC layer mechanisms for detecting transmission errors and recovering from them, and thus operate with minimal overhead.

However, it would not permit the concatenation of several small packets into one transmission frame together with selection rejection and hence retransmission on error. This has downside implications for performance, especially in AOC applications. It can thus not be readily considered until an FEC is available.

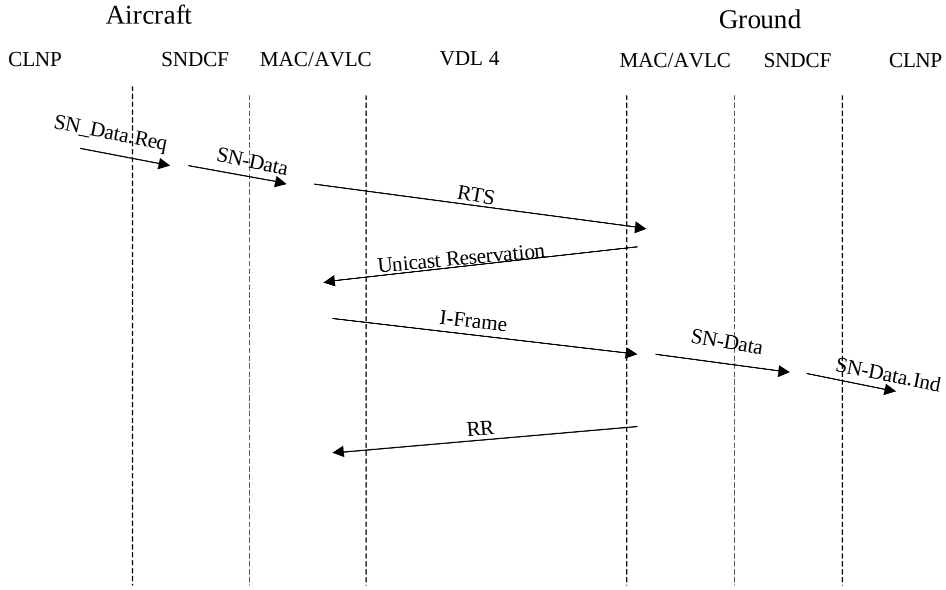

Figure 4 Data Transfer Sequence of Event

5.2.4 Conclusion

The use of the connection mode AVLC appears to be fully justified in the absence of an FEC mechanism.

However, if an FEC was available then a simplified approach appears to be possible and desirable.

5.3 Handoff

Handoff is a data link level function and should not have an impact on higher layer functions (i.e. ATN). In the original VDL Mode 2 specification (which used X.25 as a packet layer protocol over the VDL Mode 2 AVLC), Handoff is exposed to the data link user and required specific functionality to support it. The Frame Mode SNDCF has been designed to avoid this problem.

Instead, Handoff is treated by the Frame Mode SNDCF almost identically to Data Link Establishment. The Handoff Event is semantically the same as the Join Event and either precedes a Handoff or follows its successful completion. Hence a DLS exchange is either combined with the XID exchange that supports Handoff or is transferred as data over the newly established data link.

On completion of the DLS exchange, the A/GCS identifies the newly established datalink as being with the same Data link Service Provider as an existing data link. In VDL Mode 2 and by implication VDL Mode 4, there is a short period following Handoff when both old and new datalinks exist simultaneously. The compression state is copied from the old data link to the new one and the new data link is now used in preference to the old one. Typically, the old data link now drops out, although there may still be some uplinks received over it that were queued before the Handoff took place.

Note that the sequence of events for Handoff is basically the same as illustrated in Figure 3, except that the XID-CMD_HO and XID_RSP_HO formats are used. Following VDL Mode 2, it is not expected that there are any interactions with the old Ground Station. Termination of the old datalink is implicit i.e. there is no protocol exchanges between the aircraft and the old Ground Station.

5.4 Data Link Termination

Data Link Termination occurs either after a Handoff when the old data link drops out, or when the aircraft simply goes out of range of the Ground Stations operated by a given Service Provider. This is reported in each case by a “Leave” Event.

If a Leave Event occurs and there is another data link with the same Service Provider then the A/GCS notes the loss of the data link, but nothing else happens. This is simply the final stage of Handoff.

If there is no other data link with the same Service Provider then the compression state information is (optionally) saved and the network routing functions notified that the route is no longer available.

The saved compression state may be reloaded on a later Data Link Establishment with the same Service Provider, emphasising the strong similarity between Data Link Establishment and Handoff as far as the SNDCF is concerned.

The only problem here is that VDL Mode 4 inherits the data link termination conditions from VDL Mode 2. VDL Mode 2 already has a known problem in this area as the timer which the Ground Station uses to monitor loss of the AVLC connection to an aircraft is set to a very long time (of the order of 60 minutes), which is not acceptable for ATC purposes.

The VDL4 Periodic Broadcast appears to provide an ideal mechanism to provide ground based Leave Events as a lack of synchronisation bursts from an aircraft can be taken as indicating that the aircraft is out of range and hence be the trigger for a Leave Event. However, this is not described in the SARPs.

5.5 Conclusion

Using the Frame Mode SNDCF to support the integration of VDL Mode 4 into the ATN appears to be practicable and straightforward, ignoring the performance issues. The lack of an FEC is the main issue to affect how the integration is done.

Without an FEC, the emphasis is on small packets which makes it attractive to define the Join/Handoff event as occurring after data link Establishment/Handoff, and use of the connection mode AVLC.

If an FEC was available then transmission frames can be much larger, the Join/Handoff event is specified as occurring before data link Establishment/Handoff (thus leading to improved performance by combining the XID exchange with the DLS exchange), and a simpler and lower overhead data link layer can be used.

Options on way forward for ATN interface in VDL Mode 4

Issue	Option	Advantages	Disadvantage	Implications for SARPs/TM
SNDCF	Maintain ISO8208 approach	Mode 4 based on Mode 2 protocols	Number of differences between Mode 2 and Mode 4: • modified version of DLS flow control and frame formatting • link management includes some unique Mode 4 features including management of PECT by G1 counters • VSS handles reservation and access functionality • therefore no real ability to re-use Mode 2 systems – hence no system development advantages Handoff events: • introduced in Mode 2 to reduce the number of times that the compression context implemented within the SNDCF needs to be re-built as connectivity changes • not a particularly elegant method of managing connectivity changes although adding such functions to Mode 4 is relatively simple by reference to existing specifications within the ATN standards that are already used for Mode 2 Subnetwork access protocol • based on ISO8208 • contains some non-standard features in both Mode 2 and Mode 4 • so difficult to employ an off the shelf ISO 8208 solution View within ATN community • ISO 8208 no longer seen as best way of achieving subnetwork connectivity • non-connection orientated approaches now preferred • both approaches specified in Mode 3 although there are clear preferences for the connectionless approach	Few – since this is the currently specified protocol. May need to carry out an additional review of the SNDCF if this option is maintained.
	Use frame mode approach	No substantial SARPs drafting necessary to bring about introduction of frame mode Generally supported by ATN community	Some SARPs re-drafting	Straightforward: Frame mode is being specified elsewhere (current status is a working paper within ICAO) (ie no need to include it within Mode 4) Mode 4 TM will need to specify that Mode 4 supports Frame Mode and give a short specification of how the various primitives are handled within the DLS/LME and how join events XIDs are handled (solution depends on FEC issue). Within Frame Mode specification, there may e a need to introduce some special customisation. One example is to ensure that A/GCS does not concatenate frames if AVLC/no-FEC picture is maintained

Issue	Option	Advantages	Disadvantage	Implications for SARPs/TM
Introduce FEC		Allows Mode 4 to support longer messages • hence improves throughput by preventing need to segment messages	Current system tailored to ADS-B which does not need FEC (although it would obviously benefit from it) and which has little spare bits to accommodate an FEC	Bursts with invalid CRC currently rejected by MAC layer (TM 1.2.6). One approach might be to introduce a FEC for certain types of burst – specifically the compressed frame burst. The decision to reject a burst would then reside in the VSS and would define different behaviour based on the burst type. Resultant changes to TM quite simple although impact on manufacturers likely to be more significant.
AVLC	Maintain and optimise AVLC	Supports segmentation by CLNP and concatenation of short frames by providing selective reject. Such segmentation will occur frequently if no FEC is introduced	Segmentation by CLNP introduces undesirable overhead If FEC is provided, advantages of AVLC are lessened and it is left looking like a rather complex method of solving a simple problem, particularly as the VSS provides a reliable link protocol. A number of non-optimum features need to be corrected: • specify leave event based on sync bursts • clarify where re-transmission control is carried out and modify DLS procedures accordingly • allow different message length thresholds between long/short transmissions to apply to the up and downlink and in general, make sure that point to point can be made more ground controlled Specification of AVLC is not clear and, ideally, inherited text from Mode 2 should be included specifically in the TM	Significant re-drafting is required if the recommendation to import Mode 2 text is followed. This seems to be a good idea because the text as it stands at the moment is open to interpretation and really relies on experts to understand. The first stage of importing the Mode 2 text is easy but a second stage is needed in which the integrated text is reviewed and the text inevitably revised. TM changes to fix the non-optimum procedures are straightforward although clarification of the re-transmission control will need some careful consideration of procedures within the DLS.
	Replace AVLC with a simpler protocol. Two stages to consider here: • use of simple transmit acknowledge protocol • introduction of a linking protocol for long messages	Simple protocol makes it easier to implement the DLS Simple protocol probably more closely matched to VSS which already provides a reliable link service. Link protocol has a number of advantages: • it provides an alternative approach to AVLC and should be capable of cutting down on the overhead introduced by CLNP fragmentation. Hence it is worth considering even if non FEC is provided • If a FEC is provided, it may still be desirable to segment long packets in order to increase the probability of finding blocks of unreserved slots. Although such segmentation could be achieved at the CLNP it is probably better to provide segmentation at the link layer so that a) the system can support other network connections such as AOA and b) the linking overhead can be minimised A new protocol should enable priority management to be optimised	The simplified transmit acknowledge protocol is unable to transfer concatenated short frames (ie it must assign them to individual slots) and each segmented frame will generate an acknowledge – hence the system becomes less efficient if no FEC is provided for long messages – although this should be assessed against the CLNP overhead associated with segmented messages. Replacement of AVLC inevitably means a re-design of SARPs. However, the procedures needed should be much simpler than AVLC.	Significant change to TM A new procedure needs to be written although it should be relatively simple to define. Lessons can be learnt from Mode 3 A-CLDL Consideration will need to be given as to how best to optimise priority management

FEC modelling in VDL Mode 4

Prepared by Mike Shorthose, Nikos Fistas

1 Introduction

1.1 General

The VDL Mode 4 community is currently considering whether it would be beneficial to add forward error correction (FEC) to the burst for transfer of point to point messages.

This note analyses whether the addition of an FEC capability can provide significant performance enhancements for transfer of user data packets under nominal conditions (defined as bit error rate (BER) = 10^{-4}).

A simple spreadsheet model has been defined which is capable of analysing the effect of an FEC under various link conditions. The model focuses upon AVLC transfer of user data, divided into segments, using the long transmission protocol.

The effect of bit errors is to place a limit on the maximum segment length transferred across the VDL Mode 4 link. Very long segments will have a higher probability of corruption than shorter segments and hence may have to be re-transmitted. Using shorter segments can therefore minimise the number of re-transmissions necessary, but a greater number of segments are needed to transfer the message resulting in increased overhead through generation of additional request to send and clear to send transmissions. Hence, in a real system, there is a balance to be struck between these two competing factors.

It is also possible that degraded conditions (ie broadband interference, multipath, slot sharing, communication with distant aircraft etc) may occur where the BER is higher than the nominal level. In such circumstances, the addition of an FEC would be expected to have a greater beneficial impact on performance. The analysis therefore also illustrates the performance impact of an FEC on transfer of data packets under BER = 10^{-3}.

Finally, there has been discussion of whether a simpler linking protocol than AVLC could be used in order to simplify the implementation of the data link sublayer (DLS) and, potentially to reduce the format overhead required to link segments. The spreadsheet model is also used to analyse the impact in performance when an alternative link protocol (based on that used in SATCOM) is used and to compare it with the performance obtained with AVLC.

Note that all of the analysis looks at the effect of transmission loss through BER corruption only. Interference from other transmissions caused by inefficient link management (ie interference from hidden terminals and distant terminals) is not included. Hence the analysis assumes that a suitable link management concept is defined in which there is a high degree of protection for segment transfer (ie use of ground protected blocks of slots etc).

1.2 References

During its development, the modelling approach and results were discussed at two teleconferences held on 22/2/01 and 12/3/01. In these teleconferences there was participation from Eurocontrol, LFV, DFS, FAA, ADSI, SAAB and CNS Systems.

The following papers were used in the discussions and are relevant to the model:

[1] Advanced Hardware Architectures: A summary of state of the art FEC design and performance

[2] “FEC Selection Issues”, ADSI

[3] Comparison of D8PSK and GFSK modulation scheme for VDL Mode 4, AMCP WGD/11/23, VDL Mode 4 library R131, DLR/DFS

This paper presents the final version of the modelling taking into account the comments and the discussions.

2 Basic transfer detail

This section provides an overview of the basic protocol model. User data (UD) presented to the VDL Mode 4 subnetwork, is divided into one or more segments where each segment can fit into a defined number of slots on the Mode 4 link. In making this division, account is taken of:

- overhead added by VDL Mode 4 burst format
- linking overhead (ie CLNP header)
- the size of the FEC

The segments are grouped into separate flow control windows of defined size, maximum 4 segments. The final flow control window contains full length segments plus one possibly reduced size segment containing the remaining user data.

Each flow control window is transferred separately using the long transmission protocol. The stages in this transfer are:

- Stage 1: Send Request to Send (RTS)
- Stage 2: Send Clear to send (CTS)
- Stage 3: Send group of segments
- Stage 4: Send ack/CTS (modelled the same as stage 2)

Note that each flow control window is initiated by an RTS.

Individual segments/RTS/CTS/ack may be lost under the following conditions:

- Stage 1: Failure to receive RTS results in re-send by originator
- Stage 2: Failure to receive CTS results in re-transmission of the RTS followed by the CTS

- Stage 3: Will repeat until all segments are transferred. This is modelled in terms of the number of “cycles” (ie the number of times stage 3 is repeated). Note that an additional stage 2 is required for each additional cycle.
- Stage 4: As per stage 2

Other assumptions made in the model are:

- No account is taken of transmission loss due to poor slot organisation resulting in interference from hidden and distant terminals ie the model accounts for BER effects only.

3 Inputs

This section describes the inputs to the model.

Overhead associated with compressed frame burst o_1

This accounts for all fields in the compressed frame burst not containing user data. For a compressed frame burst, the following overhead is included in the burst format:

- flags (2 octets)
- reservation ID, version number, source address (4 octets)
- message ID and c/r flag (1 octet)
- link control field (1 octet)
- response reservation type (4 octets)

This total is equal to 12 octets.

Length of RTS s_1

It is assumed that the RTS occupies a single slot.

Length of CTS/acknowledgement s_2

It is assumed that the CTS/acknowledgement occupies a single slot regardless of whether an FEC is included.

Octets in first slot o_2

The length of a slot is equivalent to 32 octets of data. The last slot in a sequence should only contain 24 octets to allow 8 octets for propagation guard time. Allowing an average 0.5 octet per slot for bit stuffing, 1 slot could contain 23 octets of data including flags and reservation blocks. 2 slots could contain 54 octets, 3 slots could contain 86 octets etc.

A single slot transmission therefore contains 23 octets (including all overhead octets and flags)

Octets in subsequent slots o_3

Each additional slot can contain an average of 31.5 octets (including all overhead octets).

Segmentation overhead o_4

This represents the overhead added to the user data in order to link the segments together. For AVLC as defined in VDL Mode 4 at present, this is equal to the CLNP overhead.

This is assumed equal to 10 octets which is equivalent to the CLNP header when LREF compressed.

For other linking protocols, it is assumed that two octets are sufficient to maintain the segment sequence. A justification for this can be found by reference to the AMSS SARPs (Annex 10 Volume III Chapter 10). Each link service data unit (LSDU), which is assumed equal to the user data presented to the link layer in Mode 4, is mapped to a sequence of signal units (SUs). The linking between the SUs is accomplished by:

- a SU type identifier (2 bits)
- a sequence number (6 bits)
- a Q number (4 bits) which indicates the "precedence" of the SU (presumably used to aid priority management)
- a reference number (4 bits) which is understood to be a temporary number assigned by the link layer for the purpose of uniquely identifying the LSDU whilst it is in the process of being transferred.

The total length is therefore 2 octets. It is conceivable that only the sequence number would be required in VDL Mode 4, so the overhead may be less than 2 octets. However, 2 octets seems a reasonable budget.

Unsegmented message length o_5

This represents the unsegmented length of user data presented for transfer over the link. It is a variable used in the analysis.

BER p_1

The bit error rate is a variable input used in the analysis. Three different BER values are input into the model:

- the BER for transfer of each segment of the message, $p_{1,a}$
- the BER for transfer of the CTS, $p_{1,b}$
- the BER for transfer of the RTS. $p_{1,c}$

Depending on the investigation being carried out, BER is set either to the uncorrected BER or, if an FEC is being used, the corrected BER.

Maximum segment size on link s_3

This gives the maximum allowed segment size on the link in slots.

Flow control window N_1

This sets the maximum number of segments that can be transferred by the long transmission protocol. The maximum value is 4. A value of 1 would correspond to a simple transmit/acknowledge using the long transmission protocol.

For AVLC modelling presented in this note, a value of 4 is used.

Two alternative linking protocols are considered in the analysis: one which has the same flow control window value as AVLC and hence up to 4 segments may be transferred with a single acknowledgement. The second sets the flow control window size equal to 1 and hence results in a simple transmit/acknowledge protocol. Note that in the limit where the user data can be contained in a single segment, the two approaches become identical.

FEC model parameters

FEC performance is characterised by:

- a ***code rate*** which sets out the proportion of data bits included in a total packet of data + error correcting code. Hence a code rate of 100% includes no error correction.
- a ***code gain*** which indicates the improvement in S/N for the same BER performance after correction using an FEC.

The difficultly in using this information in a simple model is that knowledge of the BER versus S/N is required in order to assess the impact of code gain.

For the purposes of the model, it has been assumed that the performance for Mode 4 is based on the calculations made by DFS/DLR for Eurocontrol (Future VHF study) during the VDL Mode 4 validation process [3]. The results for GFSK, air/ground scenario with wide-sense-stationary uncorrelated-scattering (WSSUS) [3] are illustrated below.

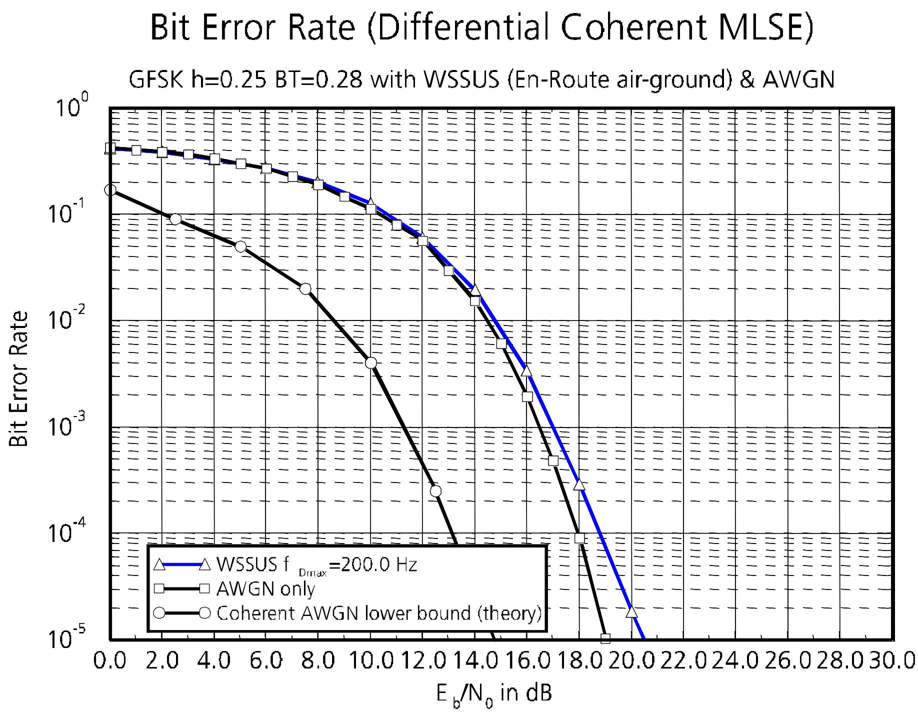

Figure 1: Bit error characteristics for en-route environment

Using this graph, it is possible to relate the corrected bit error rate to the uncorrected bit error for the same Eo/No under different assumptions of code gain. It is assumed that the impact of, for example, 1 dB code gain, is to move the WSSUS curve to the left. An uncorrected bit error rate of 10^{-3} is obtained for an Eo/No of approximately17.1 dB. Assuming that the chosen FEC introduces a code gain of 1dB, the same Eo/No produces a corrected bit error rate of approximated $3x10^{-4}$. The following table summarises the impact on BER of different assumed code gains.

Uncorrected BER	Corrected BER		
	Code gain 1dB	Code Gain 2dB	Code Gain 3dB
10^{-2}	4×10^{-3}	1.5×10^{-3}	6×10^{-4}
10^{-3}	3×10^{-4}	6×10^{-5}	3×10^{-5}
10^{-4}	3×10^{-5}	*8×10^{-6}*	*3×10^{-6}*
10^{-5}	*3×10^{-6}*	*8×10^{-7}*	*3×10^{-7}*

Figures in italics result from approximate extrapolations of the BER characteristic for uncorrected BERs below 10^{-5}.

The impact of the FEC is therefore modelled by:

- determining the assumed code gain for a particular scheme and changing the uncorrected BER to the corrected BER taken from the table above.
- inputting the code rate together with a fixed overhead to model the effect of a changed message length as a result of adding an FEC.

References [1] and [2] suggest the following two FEC schemes for evaluation:

- a block code scheme based on VDL Mode 2. This is a 97% code rate providing typically a code gain of 1dB.
- a TPC scheme providing a code gain of 2dB for a code rate of 95%. [2] provides a detailed burst proposal for this structure which is understood to involve: a) removing one octet overhead from the burst length b) adding the code octets appropriate to a 95% scheme. Note that a 2dB code gain has been chosen. [1] shows the performance of a TPC 95% code rate as providing a code gain of 2dB at an uncorrected BER of 10^{-3}. In fact, at 10^{-4} the performance is better than this. 2dB has therefore been chosen as a conservative figure.

[2] also proposes a further FEC scheme which would provide a few code bits to a single slot message such as the RTS and CTS. The code proposed is a 97% code rate and is assumed to provide about 1dB of code gain.

FEC code rate r_1

The FEC code rate gives the additional octets introduced by the FEC as a function of message length.

FEC fixed overhead o_6

This gives the additional octets introduced by the FEC scheme regardless of message size.

4 Calculations

The detailed calculations carried out in the model are described in annex A to this note.

5 Outputs

The output of the model is the number of slots required to transfer the entire user data.

6 Analysis of the impact of introducing an FEC

6.1 Output from the model

The following figures provide example outputs from the revised model. The results assume that only the data transmissions include an FEC. The RTS and CTS are sent without FEC. The analysis only considers CLNP segmentation with 10 octets of overhead.

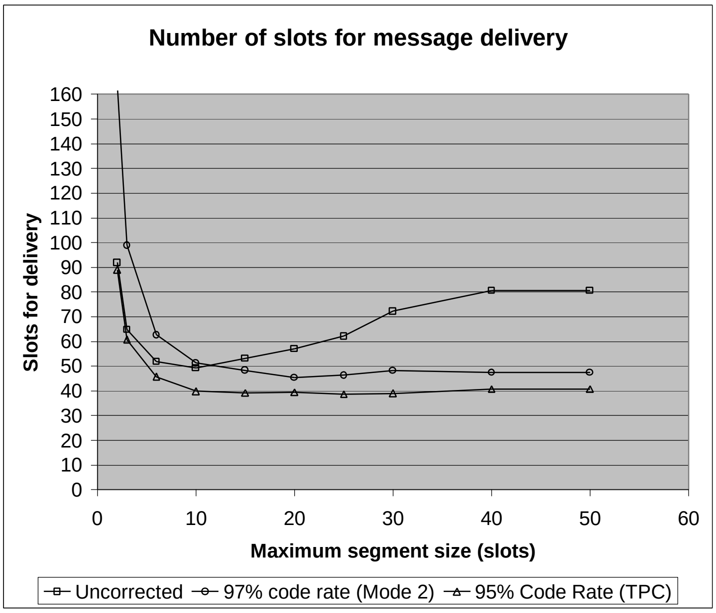

Figure 2: Unsegmented message length = 1024 octets, Uncorrected BER = 10^{4}

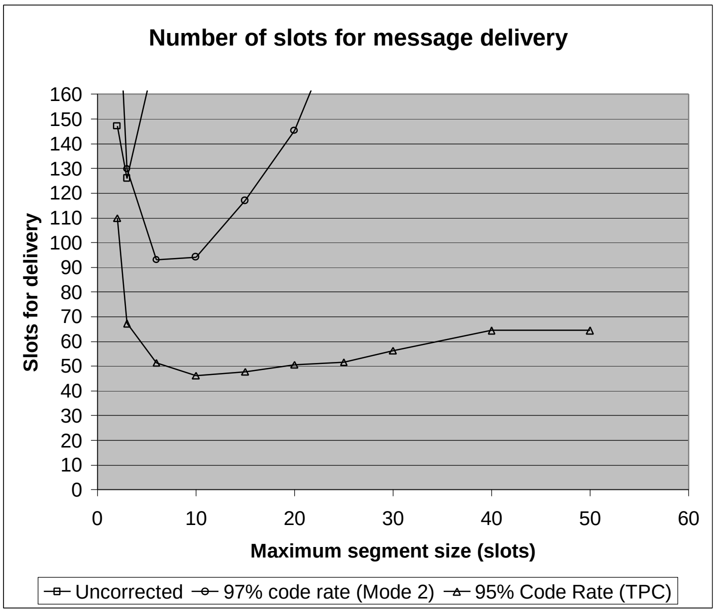

Figure 3: Unsegmented message length = 1024 octets, Uncorrected BER = 10^{3}

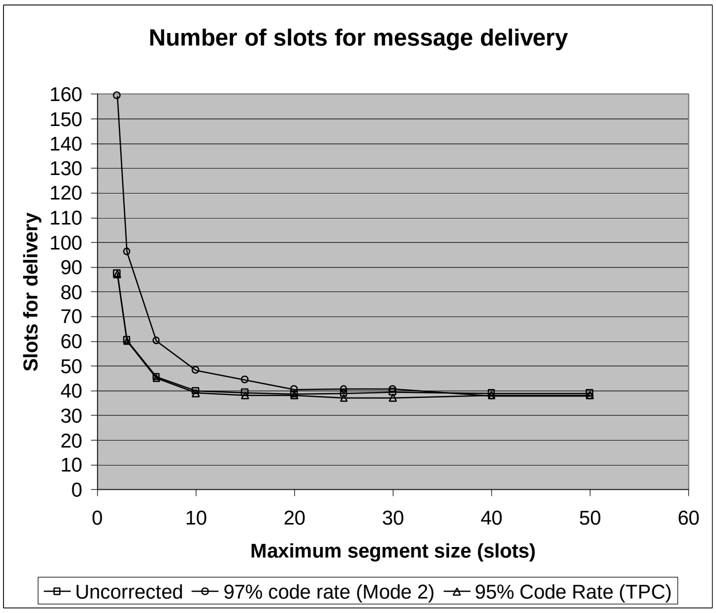

Figure 4: Unsegmented message length = 1024 octets, Uncorrected BER = 10^{-5}

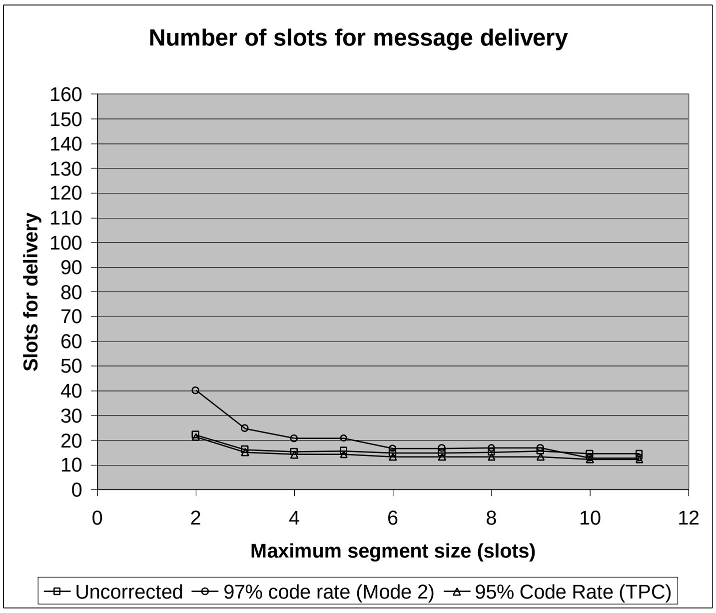

Figure 5: Unsegmented message length = 256 octets, Uncorrected BER = 10^{-4}

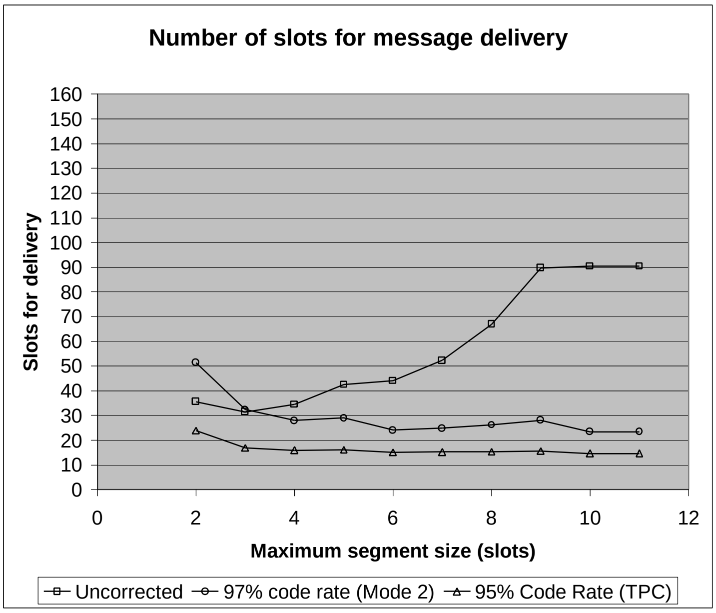

Figure 6: Unsegmented message length = 256 octets, Uncorrected BER = 10^{-3}

6.2 Summary of results

The following tables summarises the results obtained for each BER. The tables show the minimum number of slots required to transfer user data.

BER = 10-3					
User data size	no FEC	97% code rate (Mode 2)		95% code rate (TPC)	
	nb of slots	nb of slots	capacity delta	nb of slots	capacity delta
128	17,00	12,00	29,41%	9,60	43,53%
256	31,50	23,40	25,71%	14,70	53,33%
512	63,00	46,70	25,87%	25,10	60,16%
1024	126,00	93,00	26,19%	46,20	63,33%

97% code rate reduces slots required by between 25% and 29%
95% code rate reduces slots required by between 43% and 63%.

BER = 10-4					
User data size	no FEC	97% code rate (Mode 2)		95% code rate (TPC)	
	nb of slots	nb of slots	capacity delta	nb of slots	capacity delta
128	8,90	8,30	6,74%	8,20	7,87%
256	14,60	12,80	12,33%	12,30	15,75%
512	26,10	22,60	13,41%	21,50	17,62%
1024	49,20	45,30	7,93%	38,60	21,54%

97% code rate reduces slots required by between 6% and 13%
95% code rate reduces slots required by between 7% and 21%.

BER = 10-5					
User data size	no FEC	97% code rate (Mode 2)		95% code rate (TPC)	
	nb of slots	nb of slots	capacity delta	nb of slots	capacity delta
256	12,20	12,10	0,82%	12,00	1,64%
512	20,80	20,30	2,40%	21,10	**-1,44%**
1024	39,00	38,00	2,56%	37,20	4,62%

97% code rate reduces slots required by between 0% and 3%
95% code rate reduces slots required by between 0% and 4%.

In order to assess the impact of the FECs on overall channel capacity, the results have been averaged assuming that:

- 50% of messages are of length 128 octets
- 25% of messages are of length 256 octets
- 15% of messages are of length 512 octets
- 10% of messages are of length 1024 octets.

These figures are for illustrative purposes but are based on the observation that most ATSP messages are relatively short in length. A small number of longer messages results from:

- long ADS-C messages where details of a aircraft intent are included
- AOC messages
- TIS messages
- system management messages (ie IDRP derived messages)

Using this profile, the following results are obtained:

- at BER = 10^{-3}, capacity increases of 28% (97% code rate) and 50% (95% code rate) are obtained.
- at BER = 10^{-4}, capacity increases of 9% (97% code rate) and 13% (95% code rate) are obtained.
- at BER = 10^{-5}, capacity increases of 1% (97% code rate) and 1% (95% code rate) are obtained.

The results for a user data size of 1024 and a BER of 10^{-3} were repeated assuming that the CTS and RTS were also corrected using a 97% code. A small reduction in the number of slots required was obtained: 45.4 compared with 46.2 for TPC with no CTS/RTS FEC.

6.3 Conclusions

Incorporation of an FEC provides clear gains in capacity at BERs of 10^{-4} and 10^{-3} for long messages. There are also clear gains at BER if 10e-3 for short messages and some gain at 10^{-4}.

Providing an FEC for RTS and CTS does not make a significant difference to the performance.

7 Analysis of the impact of an alternative link layer protocol

7.1 Introduction

The section presents an analysis of the impact of alternative linking protocols on the efficiency of the VDL Mode 4 link.

7.2 Model results

An alternative link layer protocol to AVLC might, for example, be based on that used in AMSS. The main advantages are:

- a less complex protocol to define and implement
- a reduced linking overhead compared with AVLC.

The following table shows the impact of introducing a link protocol with 2 octets of segmentation overhead but maintaining a flow control window of 4. In this example there is no FEC.

			Total slots required for user data transfer			
	Unsegmented message length	Uncorrected BER	Uncorrected	Uncorrected + reduced segmentation overhead	FEC (95% code rate)	FEC (95% code rate) + reduced segmentation overhead
	octets		slots	slots	slots	slots
Long UD/nominal conditions	1024	10^{-4}	49.2	47.9	38.6	38.1
Long user data/degraded conditions	1024	10^{-3}	126.0	114.3	46.2	45
Short user data/degraded conditions	256	10^{-3}	31.5	27	14.7	13.9

In the absence of an FEC, the introduction of an alternative linking protocol has resulted in a reduction in the number of slots required to transfer the user data as follows:

- Long user data/nominal conditions: 49.2 to 47.9 (= 2.6%)
- Long user data/degraded conditions: 126 to 114.3 (= 9.3%)
- Short user data/degraded conditions: 31.5 to 27.0 (= 14.3%)

The reduction is greatest for larger BER and smaller user data packet size.

When an FEC is introduced, the alternative linking protocol performs only marginally better than AVLC although a performance improvement of 5% is obtained for short user data packets:

- Long user data/nominal conditions: 38.6 to 38.1 (= 1.3%)
- Long user data/degraded conditions: 46.2 to 45 (= 2.6%)
- Short user data/degraded conditions: 14.7 to 13.9 (= 5.4%)

This is because by introducing the FEC, a long segment length can be used which can contain the entire message. However, the linking protocol becomes particularly simple because all of the user data is delivered in a single segment.

7.3 Conclusions

The results show that the FEC produces the greatest reduction but that a useful further performance improvements result from reducing the segmentation overhead, particularly for shorter length messages.

8 Conclusions

The results presented in this note show that introduction of an FEC produces a significant reduction in the number of slots necessary to transfer a message.

For nominal conditions (BER = 10^{-4}), the number of slots required to transfer messages is reduced by between 7% and 22%. For a typical message load, this equates to a capacity increase of between 9% and 13%.

Under degraded conditions (BER = 10^{-3}), the transfer is difficult to sustain without a BER and a greatly increased number of slots are required to transfer the user data. The FEC has a substantial effect producing a reduction in the number of slots required to transfer the user data of between 25% and 63%. For a typical message load, this equates to a capacity increase of between 28% and 50%.Under these conditions, the performance is very close to that obtained for a BER of 10^{-4}.

Further performance benefit can be obtained by the use of an alternative linking protocol such as one based on AMSS. The primary benefits comes from a reduction in the number of octets required to support the linking of segments and through potential simplification of the DLS implementation.

In the absence of an FEC, an alternative link protocol reduces the number of slots required to transfer user data by between 3 and 15%.

In the presence of an FEC, the advantage of a simpler linking protocol is not as clear, although a reduction in the number of slots required to transfer the message by 5% is obtained for shorter message lengths.

The results quantify the self-evident advantages of using an FEC and reducing the segmentation overhead. The relevance of these results to the real system depends on the actual BER experienced on the link. VDL Mode 4 defines nominal performance as $10e^{-4}$ and it can be expected that under busy conditions, slot sharing will ensure an environment that settles around this figure. The modelling shows that introducing an FEC can provide substantial performance benefits under nominal conditions and, in addition, greatly increase the robustness of the link under degraded conditions.

The case for a simpler link protocol is less clear and needs further discussions. There are some clear performance benefits as well as the possibility to provide a standard that is more straightforward to implement. However, that needs to be set against the difficulty at this stage of introducing a new DLS

Appendix A: Details of calculations

This section provides details of the model used to model the transfer performance. It builds on the inputs defined in section 3.

Size of FEC o_{12}

This is the size of the FEC for a full segment and includes the additional octets generated as a result of the FEC code rate and also the fixed additional overhead.

$$o_{12} = (1 - r_1)((s_3 - 1) \times o_3 + o_2) + o_6$$

Maximum segment size (user data including header) o_7

$$o_7 = (s_3 - 1) \times o_3 + o_2 - o_1 - o_{12}$$

This is the maximum size of the user data + segmentation overhead in a segment.

Maximum segment size (user data only) o_8

$$o_8 = o_7 - o_4$$

This is the maximum size of the user data in a segment.

Number of segments of maximum size N_2

$$N_2 = \mathrm{INT}\left(\frac{o_5}{o_8}\right)$$

This is the number of full length segments that are required to transfer the user data

Number of segments of maximum size (corrected) N_3

In the event that the maximum segment size (user data only) divides exactly into the unsegmented message length, for the purposes of later calculations, the number of segments of maximum size is corrected by subtracting 1. This ensures that there is always a remainder segment even if it is a full size. (note this is done to make the later stages of the simulation calculation easier – it is an implementation artefact)

$$\begin{aligned}
&\text{IF} \quad (o_5 - o_8 \times N_2) = 0 \quad \text{THEN} \\
&\qquad N_3 = N_2 - N_1 \\
&\text{ELSE} \\
&\qquad N_3 = N_2 \\
&\text{END IF}
\end{aligned}$$

Size of remainder (user data) o_9

$$o_9 = o_5 - o_8 \times N_3$$

This calculates the user data in the final (remainder) segment

Size of FEC (remainder) o_{13}

This is the size of the FEC for the remainder segment and includes the additional octets generated as a result of the FEC code rate and also the fixed additional overhead.

$$o_{13} = \left((1 - r_1)(o_9 + o_1 + o_4) + o_6\right) / r_1$$

Size of remainder (user data + o/h)) o_{10}

$$o_{10} = o_9 + o_1 + o_4 + o_{13}$$

This calculates the user data plus overhead (compressed frame burst, segmentation, FEC) in the final (remainder) segment

Slots required for remainder s_4

$$\begin{array}{l} \text{IF} \quad o_{10} > o_2 \quad \text{THEN} \\ \qquad s_4 = 1 + \text{INT}\left((o_{10} - o_2)/o_3\right) \\ \text{ELSE} \\ \qquad s_4 = 1 \\ \text{END IF} \end{array}$$

This is the number of slots required to transfer the final segment.

Total number of segments N_4

$$N_4 = N_3 + 1$$

This indicates the total number of segments needed to transfer the user data packet

Number of full length flow control windows m_1

$$\begin{array}{l} \text{IF} \quad \dfrac{N_4}{N_1} = \text{INT}\left(\dfrac{N_4}{N_1}\right) \quad \text{THEN} \\ \qquad m_1 = \dfrac{N_4}{N_1} - 1 \\ \text{ELSE} \quad m_1 = \text{INT}\left(\dfrac{N_4}{N_1}\right) \\ \text{END IF} \end{array}$$

Indicates how many groups containing the maximum number of full length segments need to be transferred.

Number of segments in final flow control window N_5

$$\text{IF} \quad \frac{N_4}{N_1} = \text{INT}\left(\frac{N_4}{N_1}\right) \quad \text{THEN}$$

$$N_5 = N_1$$

$$\text{ELSE} \quad N_5 = N_4 - N_1 \times \text{INT}\left(\frac{N_4}{N_1}\right)$$

$$\text{END IF}$$

Indicates how many segments are transferred in the final group.

Total full length segment size o_{11}

$$o_{11} = o_7 + o_1 + o_6$$

This is the total number of octets in a segment including all overhead.

Segment delivery probability p_2

$$p_2 = \left(1 - p_{1,a}\right)^{(8o_{11})}$$

This is the probability of successfully delivering a segment taking account of the bit error rate and FEC correction. Note that the FEC model is implemented by calculating a corrected BER. Therefore, the segment delivery probability is calculated assuming that all bits must be delivered uncorrupted.

RTS delivery probability p_3

$$p_3 = \left(1 - p_{1,c}\right)^{(8o_2)}$$

As per the segment delivery probability only using the RTS length equal to the maximum number of octets contained in a single slot.

CTS delivery probability p_4

$$p_4 = \left(1 - p_{1,b}\right)^{(8o_2)}$$

As per the segment delivery probability only using the CTS length equal to the maximum number of octets contained in a single slot.

Final segment delivery probability p_5

$$p_5 = \left(1 - p_{1,a}\right)^{(8o_{10})}$$

As per the segment delivery probability only using the length of the final segment.

Stage 1: Expected number of RTS transmissions N_6

$$N_6 = \frac{1}{p_3}$$

Taking account of RTS delivery probability, this gives the average number of transmissions necessary to delivery the RTS. It includes the effect of re-transmissions.

Stage 2: probability density function $p_6(n)$

Stage 2 consists of sending the CTS. IF the CTS fails, then a new RTS is issued by the initiating station. Successful delivery of the CTS then occurs if both the RTS and the CTS are successful. The probability density function for the number of times the CTS must be sent is given by:

$$p_6(n) = p_4$$
$$p_6(n) = (1 - p_4)p_4 p_3 (1 - p_4 p_3)^{n-2} \quad \text{for} \quad n > 1$$

Stage 2: average number of cycles N_7

$$N_7 = \sum_{n=1}^{10} n p_6(n)$$

This is the average number of times stage 2 must be repeated taking account of transmissions delivery probability.

Stage 2: CTS number $N_8(n)$

The number of CTS transmissions necessary to complete stage 2 depends on a) whether the RTS fails b) whether, given a successful RTS, the CTS then fails. The number of CTS transmissions as a function of the number of stage cycles is given by:

$$N_8(1) = 1$$
$$N_8(n) = 2 + \frac{(n-2)p_3(1 - p_4)}{(1 - p_3 p_4)} \quad \text{for} \quad n > 1$$

Stage 2: average number of CTS N_9

$$N_9 = \sum_{n=1}^{10} N_8(n) p_6(n)$$

This is the average number of times the CTS must be re-sent taking account of transmission delivery probability.

Stage 2: RTS number $N_{10}(n)$

The number of RTS transmissions necessary to complete stage 2 depends on a) whether the RTS fails b) whether, given a successful RTS, the CTS then fails. The number of RTS transmissions as a function of the number of stage cycles is given by:

$$N_{10}(n) = n - 1$$

Stage 2: average number of RTS N_{11}

$$N_{11} = \sum_{n=1}^{10} N_{10}(n) p_6(n)$$

This is the average number of times the RTS must be re-sent taking account of transmission delivery probability.

Stage 3: Expected number of segment transmissions N_{12}

$$N_{12} = \frac{N_1}{p_2}$$

Taking account of segment delivery probability, this gives the average number of transmissions necessary to deliver each segment multiplied by the number in each flow control window. It includes the effect of re-transmissions.

Stage 3: Number of phase 3 repeats PDF $p_7(n)$

$$p_7(1) = p_2^{N_1}$$
$$p_7(n) = \left(1 - (1 - p_2)^n\right)^{N_1} - \left(1 - (1 - p_2)^{n-1}\right)^{N_1}$$

The full flow control window will be delivered when each segment has been delivered. Hence, the cumulative probability of completing the flow control window is the product of the cumulative probability for each segment. The PDF is just the difference between the cumulative probability for n repeats and that for n-1 repeats.

Stage 3: Expected number of cycles N_{13}

$$N_{13} = \sum_{n=1}^{10} np_7(n)$$

This is the average number of times a group of segments must be sent to complete stage 3 taking account of transmission delivery probability.

Final segment group

This repeats the calculations set out above, taking account of a) the possibility that there is a reduced number of segments in the final flow control window b) the possibility that the final segment is shorter.

The calculated quantities are:

- expected number of segment transmissions (full length) N_{14}
- expected number of segment transmissions (remainder) N_{15}
- expected number of final group repeats PDF $p_8(n)$
- average number of final group cycles N_{16}

Summary of full transfer

This part brings together the calculations to derive a total number of slots to deliver the user data packet.

Firstly, assuming no loss:

A first stage is required for each flow control window N_{17}

$$N_{17} = m_1 + 1$$

Two second stages (the CTS and the acknowledge) are required for each full length flow control window. N_{18}

$$N_{18} = 2m_1$$

A third stage is required for each full length flow control window N_{19}

$$N_{19} = m_1$$

A second stage (final group) is required for the final flow control window N_{20}

$$N_{20} = 2$$

A third stage (final group) is required for the final flow control window. N_{21}

$$N_{21} = 1$$

Next assuming that there is a loss of transmissions, additional stage 2's are required slotting in between each additional transmission of the flow control window.

Number (with loss)

The total second stages (the CTS and the acknowledge) is given by: N_{22}

$$N_{22} = N_{18} + (N_{13} - 1)m_1$$

The total second stages (final group) required for the final flow control window is given by: N_{23}

$$N_{23} = N_{20} + N_{16} - 1$$

RTS transmissions occur in stage 1 and also stage 2. The total expected numbers are given by:

First stage N_{24}

$$N_{24} = N_{17} N_6$$

Second stage N_{25}

$$N_{25} = N_{22} N_{11}$$

Second stage (final group) N_{26}

$$N_{26} = N_{23} N_{11}$$

The total slots required to accommodate these RTS transmission is given by: s_5

$$s_5 = (N_{24} + N_{25} + N_{26})s_1$$

CTS transmissions occur in stage 2. The total expected numbers are given by:

Second stage N_{27}

$$N_{27} = N_{22} N_9$$

Second stage (final group) N_{28}

$$N_{28} = N_{23} N_9$$

The total slots required to accommodate these CTS transmission is given by: s_6

$$s_6 = (N_{27} + N_{28}) s_2$$

Full length segment transmissions occur in stage 3. The total expected numbers are given by:

Third stage N_{29}

$$N_{29} = N_{19} N_{12}$$

Third stage (final group) N_{30}

$$N_{30} = N_{21} N_{14}$$

The total slots required to accommodate these CTS transmission is given by: s_7

$$s_7 = (N_{29} + N_{30}) s_3$$

Reduced length segment transmissions occur in stage 3 final stage only. The total expected numbers are given by:

Third stage

Third stage (final group) N_{31}

$$N_{31} = N_{21} N_{15}$$

The total slots required to accommodate these CTS transmission is given by: s_8

$$s_8 = N_{31} s_4$$

Output

Total slots required for user data transfer s_9

This is given by: $s_9 = s_8 + s_7 + s_6 + s_5$

FEC SELECTION ISSUES

by Steve Friedman

February 27, 2001

ADSI

ADSI, Inc.
7900 Wisconsin Avenue, Suite 201
Bethesda, Maryland 20814
+1-301-652-5306
adsi@adsi-m4.com

CONTENTS

1. BACKGROUND

Because of concerns regarding the system performance of VDL/4 data transfer, the VSG is determining whether the addition of FEC can improve the overall system capacity. As part of that effort, I was tasked with the action item of providing a background paper on available codes and the implications of their use. This paper attempts to answer those initial questions along with raising issues regarding the implementation of an FEC. The primary purposes for this paper are:

- to provide sufficient details for Mike Shorthose's analysis of the performance of the link given an FEC characteristic,
- to highlight for the committee the potential performance of modern FECs, and
- to highlight for the committee members some of the design work remaining before the completion of the FEC can be finalized.

2. FEC CHARACTERISTICS

These are some basic issues in the FEC selection.

2.1 Systemic vs Non-systemic

A systemic code has separate data bits and code bits. This enables a cheaper implementation: simply ignore the code bits and pick out the data bits. A non-systemic code would require all implementations to implement the FEC decode.

2.2 Soft-decision Decoding vs. Hard-decision Decoding

Soft-decision decoding means that the demodulator outputs a measure of goodness along with each decoded bit. This enables the decoder to concentrate more of its power on the bits that the demodulator suspects as being in error. Even a few bits of quality metric can enable 2 or 3 dB of coding gain. The availability of a soft decision algorithm greatly improves the performance of the code.

2.3 Block Code vs. Trellis Code

A trellis code is easier to use with variable length transmissions, but these codes suffer from reduced performance to burst errors. Trellis code performance ends up devolving to the interleaver performance used to spread out burst errors. Block codes offer superior performance to burst errors, but are harder to adapt to variable length transmissions. However, because of the slot structure of VDL/4, it may be easier to adapt a block code as there is no channel capacity loss for transmitting an entire slot rather than a partial slot. Finally, trellis codes do not work with lightweight codes. Consequently, I presume that we will select a block code and modify the framing structure to suit the code.

2.4 Random Bit Error vs. Burst Error

Isolated errors (typical in an AWGN environment) that place equal stress on all parts of the code are easier to correct than burst errors that place all of the stress on a small part of the code.

3. FEC OPTIONS

Because we must choose a single code and use it throughout the flight regime, we cannot approach Shannon's limit: we will tend to use either too much or too little code. Modern codes get very close to Shannon's limit with long transmissions. However, the power of the code depends on the interleaver and the framing structure of the transmission.

3.1 VDL/2 – Lightweight Reed-Solomon Code

The VDL/2 FEC is a very lightweight code (a rate of 97%) and provides about 1 dB of coding gain with hard-decision decoding. (During the validation flight trials of circularly polarized GNSS uplinks, the FEC was correcting at least half of all received blocks – a far greater effect than what we envisioned when we designed the code.) When the code was selected, there was active research in developing soft-decision decoding algorithms for Reed-Solomon codes. Recently a number of superior codes have been developed (some that reach within 0.0045 dB of Shannon's limit – the so-called Low-Density Parity-Check (LDPC) Codes). Consequently, work in soft-decision Reed-Solomon codes has died down.

3.2 VDL/3

Since VDL/3 FEC is not state-of-the-art, I did not do the research on its latest specification or performance and it is not considered here.

3.3 Turbo Codes

Turbo-codes, concatenated code with an iterative soft-decision input soft-decision output decoding process, are currently the optimal known codes in active use today. That is, using soft-decision decoding, the decoder sends the received data through one decoder, then the second decoder, back through the first decoder, etc. Since a code is correcting the mistakes that leaked through the other code, two very simple codes combine to form an extremely powerful code. This iterative process takes time, but data blocks do not contain reservations so there is no concern with delayed reservations. These codes either use to block codes (and are called Turbo Product Codes (TPC)) or two trellis codes (and are called Turbo Convolutional Codes (TCC)).

Please see the attached presentation from AHA regarding performance of turbo product codes (TPC). As can be seen on pages 53 and 67, the performance of even a very lightweight code is rather impressive.

3.4 LDPC Codes

These codes promise to be about 1 dB better than turbo codes ; however, more research is needed to reduce the operational complexity of these codes. A few years from now this code family would merit serious consideration.

4. ANCILLARY ISSUES

4.1 Fixed Transmission Protection

Besides data transfer, the other long transmissions that are immediately envisioned are long fixed ground transmissions. Uplinks such as weather, traffic, DOS, GSIFs, can benefit from trading off link overhead for FEC to improve the reception probability. However, these transmissions include reservation data. As a community, we need to decide whether we want to allow FEC on reservation-containing bursts.

4.2 Transmission Structure

4.2.1 Unique Word

Because a transmission containing FEC must be differentiated from a transmission that does not contain FEC, we must use a different unique word to enable an immediate differentiation of the coding structure.

4.2.2 Framing Structure

The current framing structure (ISO 3309) adds 8 bits + 8 bits per frame of overhead for the flags, 1.6% (on average) overhead for the bit stuffing, and 16 bits per frame overhead for the CRC. Even if the selected FEC could provide a BER of 10^-19 (which some turbo block codes are capable of), some CRC would still be required. Perhaps not reducing the size of the CRC will simplify acceptance of the data link. Regardless, if we can recover the bit stuffing overhead, then we can reduce the cost of the FEC. And, since the bit stuffing is a random process, removing it allows the VSS to know the exact number of slots required (without having to fudge for the unknown amount of bit stuffing required).

There is no need to shorten any transmission less than the standard end of a slot – there is no reduction in performance to not transmit FEC bits to complete a slot. On the other hand, the first bit of FEC beyond the end of the slot boundary is a very expensive bit.

4.2.3 NRZI Encoding / Scrambling

There is a resynchronization issue with NRZI encoding: because a high tone indicates not a one or zero, but rather a change from the previous bit, all errors become double-bit errors as the subsequent bit is interpreted as a change in the transmitted frequency. Consequently, if a symbol is interpreted incorrectly, then two bits (rather than one) will be decoded incorrectly. NRZI was used because, with HDLC bit stuffing, it eliminated the need for scrambling. However, now that we are including FEC, we should not add to its burdens. Further, if the replaces ISO 3309 framing (and its bit insertion) with a block structure, as recommended, NRZI is not sufficient to eliminate the need for a scrambler.

4.2.4 Proposed Transmission Structure

I propose a structure that has less overhead than the existing ISO3309 structure. This will reduce the downside of adding FEC. Assuming that a lightweight code (that is, a block code) will be selected, we will need to transmit the length of the message in order to determine the block size.

4.2.4.1 Unique Word

To retain the existing air interface for the already validated sync bursts, we need a mechanism to ensure that a receiver can differentiate between transmissions that contain FEC and those that do not. The standard method is to use different unique words. Retaining the 24-bit length of the existing preamble, a 24-bit pattern can be chosen based on its good auto-correlation pattern by evaluating all 2^24 patterns and selecting the best pattern.

4.2.4.2 Encoding

A low tone will be a zero. A high tone will be a one. We should not use NRZI encoding because it causes double bit errors.

4.2.4.3 Scrambling

Deleting bit insertion and NRZI encoding means that scrambling is now necessary.

4.2.4.4 Transmission Length

Because there is no benefit to the system to not utilize an entire slot, we need only transmit the number of slots in the current transmission, thus reducing the dynamic range on this element. Since the loss of this field means the loss of the entire transmission, and because short blocks need comparatively more code to be transmitted successfully, this field will need its own code. Since very short transmissions will not be protected by FEC, there is no reason to define a short TL field for very short transmissions and a longer TL field for the remaining transmissions. The committee will determine the number of bits for the maximum length and the number of bits for the code. (As an example, using a rate ½ code and 8 bits for the length yields a maximum transmission of 8 KB and a 16-bit physical layer header.)

The net result is an additional byte in the physical layer header as the initial flag of the burst format is replaced with a 2 byte physical layer header.

4.2.4.5 FEC

As the data is a single block with codes in both the columnar and row direction, the readout is helical (the interleaver is implicit in the readout). One or more families of TPC codes (appropriately shortened) is used to provide codes at the various lengths required.

4.2.4.6 Burst

A transmission will consist of multiple bursts, all protected by a single FEC that is defined by the transmission length. If the amount of data to be sent is less than that required for the code, the transmitter will add the necessary number of zero bytes. This burst is basically as described in Figure 1-5 of the Technical Manual, except that the flags and bit stuffing are not used.

The net result is a savings of one octet per slot of the transmission: delete the bit stuffing (which is accounted for as 1 byte per slot or 3% overhead) plus the final flag and replace with a burst length field. If a TPC with 0.95% code rate were selected, this would add only 2% to the length of the transmission.

4.2.4.6.1 Burst Length

This is the length of the burst in bits. It is a 10-bit field (maximum burst length is 1KB). To ensure that the data is octet aligned, we either end up with 6 spare bits or we steal two bits from the first octet (available bits include the version field and the TCP change bit). Of course, a 1 KB burst is nearly 30 slots. If the maximum burst (not transmission length) is 512 bytes, then a 9-bit field would suffice and we could accommodate it by stealing the TCP change bit.

4.2.4.6.2 Data

The rest of the data is unaffected by this change.

4.2.4.6.3 CRC

To ensure that the system correctly breaks the transmission into the correct bursts, we retain a CRC. To simplify validation and acceptance of the protocol, we retain the existing 16-bit CRC.

4.3 CTS Encoding

Material needs to be added in the CTS encoding regarding the additional slot(s) required for the FEC. The current Technical Manual does not provide guidance on the amount of bit stuffing to reserve; which would otherwise need to be rectified.

Change Proposal: FEC

Title: Support for FEC
Author: Steve Friedman
Date Submitted: 23MAR01

Rationale: The VSG, upon investigation of the performance of the GFSK physical layer for long transmissions typical of data communications has determined that FEC will enhance the capacity of the link at bit error rates of 10^-3 to 10^-4 with minimal degradation at higher bit error rates.

SARPs Changes (changed from AMCP/7-DP/2, Appendix A to the Report on Agenda Item 1):

6.9.5.2 PROTOCOL DEFINITION FOR GFSK

6.9.5.2.1 *Modulation scheme.* The modulation scheme shall be GFSK. ~~The first bit transmitted (in the training sequence) shall be a high tone and the transmitted tone shall be toggled before transmitting a 0 (i.e., non-return to zero inverted encoding).~~

6.9.5.2.2 *Modulation rate.* Binary ones and binary zeros shall be generated with a modulation index of 0.25 +- 0.03 and a BT product of 0.28 +- 0.03, producing data transmission at a bit rate of 19.200 bits/sec +- 50 ppm.

6.9.5.2.3 ~~TRAINING~~ *Transmission sequence.* A transmission, with or without FEC, shall consist of the following segments

6.9.5.2.3.1 *Transmitter power stabilization.* The first segment of the training sequence is the transmitter power stabilization, which shall have a duration of 16 symbol periods. The transmitter power level shall be no less than 90% of the steady state power level at the end of the transmitter power stabilization segment.

6.9.5.2.3.2 *Synchronization and ambiguity resolution.* The second segment of the training sequence shall be ~~the~~ a 24-bit binary sequence defined below ~~0101 0101 0101 0101 0101 0101,~~ transmitted from left to right immediately before the start of the data segment.

6.9.5.2.3.3 *Data transmission.* The transmission of the first bit of data shall start 40 bit intervals (approximately 2083.3 microseconds) +- 1 microsecond after the nominal start of transmission.

6.9.5.2.3.4 *Transmission decay.* The transmitted power level shall decay at least 20 dB within 300 microseconds after completing a transmission. The transmitter power level shall be less than –90 dBm within 832 microseconds after completing a transmission.

6.9.5.2.3.5 *Propagation guard time.* The propagation guard time including receiver to transmitter turnaround shall be 1 250 microseconds.

6.9.5.2.4 Definition for no error correcting scheme. Transmissions that do not include FEC shall conform to Sections 6.9.5.2.1, 6.9.5.2.2, 6.9.5.2.3, and 6.9.5.2.4.

6.9.5.2.4.1 *Bit encoding.* The first bit transmitted (in the training sequence) shall be a high tone and the transmitted tone shall be toggled before transmitting a 0 (i.e., non return to zero inverted encoding).

6.9.5.2.4.2 *Synchronization and ambiguity sequence.* The 24-bit binary sequence shall be 01010101 0101 0101 0101 0101.

6.9.5.2.5 *Definition for error correcting scheme.* Transmissions that include FEC shall conform to Sections 6.9.5.2.1, 6.9.5.2.2, 6.9.5.2.3, and 6.9.5.2.5.

6.9.5.2.5.1 *Bit encoding.* A binary one shall be transmitted as a high tone and a binary zero shall be transmitted as a low tone.

6.9.5.2.5.2 *Synchronization and ambiguity sequence.* The 24-bit binary sequence for short transmissions (less than seven slots) shall be TBD. The 24-bit binary sequence for long transmissions shall be TBD.

6.9.5.2.5.3 *Bit scrambling.* Under Mode 4 operation with forward error correction, bit scrambling, as specified in Section 6.4.4.1.4 shall be performed on each burst, starting after the synchronization sequence. The scrambling sequence shall be reinitialized on each burst effectively providing a constant overlay for each of the Mode 4 fixed length bursts.

6.9.5.2.5.4 *Physical layer header.* The physical layer header for a short transmission shall be the sequence $n_0\ n_1\ n_2\ p_0\ p_1\ p_2\ p_3\ p_4$, transmitted from left to right The physical layer header for a long transmission shall be the sequence $n_0\ n_1\ n_2\ n_3\ n_4\ n_5\ n_6\ n_7\ p_0\ p_1\ p_2\ p_3\ p_4\ p_5\ p_6\ p_7$, transmitted from left to right. The n bits shall be the length of the transmission in slots and the p bits shall be the parity bits generated using the following equation and the appropriate generator matrix:

$[n_0 \ldots n_{max}]\ H^T = [p_0 \ldots p_{max}]$

H_{short}=[]
H_{long}=[]

6.9.5.2.5.5 *Forward error correction code.* TBD.
[Author's Note: We can either select the same 97% code for both long and short transmissions, or select one code for transmissions too short to support TPC (i.e., less than 3 slots in length), and a TPC code for longer transmissions. The simplest specification is for a single code of a single slot (with erasures for the first, partial slot) and interleave across all of the slots in the transmission. A single-error correcting BCH code of (255,247) exists. The next easiest specification is for a family of 97% codes, one for each transmission length up to 8 or 16 slots (with optional interleaving across the blocks). Finally, with more work, we can identify a series of 97% or 95% codes for varying burst lengths.]

6.9.5.2.2.6 *Interleaving.* TBD (if needed)
6.9.5.3.4.2 *Measurement of transmission length.* When the non-FEC training sequence has been detected, the channel busy state shall be held for a period of time at least equal to 5 milliseconds, and subsequently allowed to transition to the idle state based on measurement of channel power. When an FEC training sequence has been detected, the channel busy state shall be held for a period of time equal to the number of indicated slots.
Note. To protect against a transmission extending into the slot subsequent slot (e.g., because of propagation delay or clock skew), the receiver must use time and not the number of slots.

Technical Manual changes (from TM_AMCP_rev1 chv21)
1.1
1.2
1.3
1.3.1
1.3.2 **Burst format**

VSS bursts that do not contain FEC shall conform to ISO 3309 frame structure except as specified in Table 1-5a. VSS bursts that contain FEC shall conform to Table 1-5b or Table 1-5c; the burst length shall be the number of bytes in the frame including the CRC. The maximum burst length shall be N1 bits.

Note 1.— A burst occupying a single slot has 24 octets for data. Thus, assigning 8 bits for bit stuffing and 2 octets for the flags (in the case of a non-FEC burst) OR 8 bits for the physical layer header, 8 bits for the block length header and 8 bits for the short FEC (in the case of an FEC burst), a maximum single-slot burst has a value of 'n' equal to 21. Note that a burst can be up to N1 bits in length and can therefore occupy more than one slot. Non-FEC bBursts can consist of the single block of data between two flags, as illustrated in Table 1-5a, or can consist of a number of blocks of data with each block separated from the next by a flag. FEC bursts can consist of a single block of data as illustrated in Table 1-5b, or can consist of a number of blocks of data as indicated by the individual burst lengths.

Table 1-5a. Non-FEC (ISO 3309 based) burst format

Description	**Octet**	**Bit number**							
		8	7	6	5	4	3	2	1
flag	-	0	1	1	1	1	1	1	0
See Table 1-5d for the definition of the data									
flag	-	0	1	1	1	1	1	1	0

Table 1-5b. FEC short block (specified length) burst format

Description	**Octet**	**Bit number**							
		8	7	6	5	4	3	2	1
Burst length flag, burst length	1	0	l_6	l_5	l_4	l_3	l_2	l_1	l_0
See Table 1-5d for the definition of the data									

Table 1-5c. FEC long block (specified length) burst format

Description	**Octet**	**Bit number**							
		8	7	6	5	4	3	2	1
Burst length flag, burst length	1	1	L_{14}	l_{13}	l_{12}	l_{11}	l_{10}	l_9	l_8
	2	l_7	l_6	l_5	l_4	l_3	l_2	l_1	l_0
See Table 1-5d for the definition of the data									

Table 1-5d. Burst format

Description	Octet	Bit number							
		8	7	6	5	4	3	2	1
reservation ID (rid), version number (ver)	1	s_{27}	s_{26}	s_{25}	ver_3	ver_2	ver_1	rid	1
source address (s)	2	s_{24}	s_{23}	s_{22}	s_{21}	s_{20}	s_{19}	s_{18}	s_{17}
	3	s_{16}	s_{15}	s_{14}	s_{13}	s_{12}	s_{11}	s_{10}	s_9
	4	s_8	s_7	s_6	s_5	s_4	s_3	s_2	s_1
message ID (mi)	5	in_k	mi_k	. . . . . .		mi_3	mi_2	mi_1	mi_4
information	6								
	7 - n-5			. . . .					
	n-4								
reservation data (rd)	n-3		in_1	rd_k	. . . . . .				
extended reservation ID (erid)	n-2	$erid_k$	. . . . .			$erid_1$			rd_1
CRC (c)	n-1	c_9	c_{10}	c_{11}	c_{12}	c_{13}	c_{14}	c_{15}	c_{16}
	n	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8

. . . . Denotes variable length field

Note 2.— The least significant bit of the first octet is a one in a burst). Note also that ATN applications are supported by compressed frame bursts (see Section 1.4.2.2.2).
[Author's Note: This note was stricken when bit 1 was changed to the TCP change flag – has this been accepted / adopted or is it still just proposed]

Note 3.— The lengths of the message ID and information fields depend on the message type being used.

Note 4.— The lengths of the extended reservation ID and reservation fields depend on the reservation protocol being used. The reservation field may also be subdivided into further fields, depending on the reservation protocol.

Note 5. – For the specified length burst formats, the minimum normal length is 3 bytes – 1 byte of data and 2 bytes of CRC. A burst length of zero is a pad byte that should not be passed to a higher layer. A burst length of 1 is impossible and a burst length of 2 is useless.

SECTRA

FEC in VDL Mode 4

Company Confidential Document Make it simple...

Contents

1 Executive Summary

The VDL Mode 4 standard is in many ways an excellent standard. In order to further enhance robustness and reliability for future applications (e.g. Com-applications) we suggest some modifications at the physical and link level.

Sectra Wireless Technologies AB and C.N.S. Systems AB jointly suggest the changes described below in chapter 3 to the VDL Mode 4 standard.

The suggestion herein is inspired by the work presented by ADSI where Forward Error Correction (FEC) is introduced into the VDL Mode 4 data link.

A main goal of the suggested solution is to adapt the ADSI ideas to the slotted structure of the VDL Mode 4 standard. By for instance removing the HDLC framing (as suggested by ADSI) bits are made available for a powerful FEC scheme.

According to ideas presented by M. Shorthose during the teleconference of March 12th, 2001, and the following discussion this suggestion will handle multislots of length up to 10 slots.

By applying Forward Error Correction on each individual slot this suggestion offers low implementation complexity for any message length and powerful random and burst error correction/detection. This will also result in negligible processing delay due to FEC.

The FEC scheme suggested herein is based solely on non-patended features.

2 Reading Instructions

Chapter 3 describes the suggested changes.

2.1 Document History

Revision	Changes	Date	Responsible
1.0	First issue	2001-03-30	FQ

2.2 Reference Documents

Standards

[1] ”VDL Mode 4 SARPs”, ICAO/AMCP.

[2] ”Manual on Detailed Technical Specifications for the VDL Mode 4 Data Link”, ICAO/AMCP.

2.3 Abbreviations

Abbreviation	Meaning
CRC	Cyclic Redundancy Check
FEC	Forward Error Correction
GF	Galois Field
LI	Length Indicator
MER	Message Error Rate
RFU	Reserved for Future Use
RS	Reed-Solomon
SNR	Signal-to-Noise Ratio

3 Suggested Changes

3.1 Transmitter Power Stabilization (Ramp-Up) and Decay

The time period for transmitter power stabilization should be 8 bits long (~ 416 microseconds).

The time period for transmitter power decay should be 8 bits long (~ 416 microseconds).

During these periods the transmitter should send the unmodulated center frequency.

Ramp up/down should be according to a raised cosine function.

Note: in the current VDL4 standard ramp-up takes 16 bits and decay 16 bits.

3.1.1 Rationale

- The suggested time period is long enough to allow for easy implementation while fulfilling the spectrum mask requirement.
- Sending the center frequency means maximizing the spectral distance to the adjacent channels (i.e. minimizing adjacent-channel interference).
- The raised cosine function minimizes the spectrum growth due to power switching.

3.2 Synchronization Sequence

The synchronization sequence SS should be 16 bits (~ 833 microseconds) long.

The SS should have good autocorrelation properties, e.g. as the following one:

SS = 0001 1101 1101 0010 (hex 1DD2)

Note: this sequence replaces the HDLC flag of the current VDL4 standard.

3.2.1 Rationale

The suggested synchronization sequences has good autocorrelation. It enables synchronization with high probability also on noisy channels.

3.3 Data Encoding

NRZ-L encoding of the data stream should be applied[1].

3.3.1 Rationale

NRZ-L does not introduce error propagation.

3.4 Scrambling

Scrambling of the data segments including LIs and FEC should be applied.

In single and half slots the first part of the sequence generated by the maximum-length polynomial $1 + x^4 + x^9$ should be used. The sequence should be tapped from the first (i.e. LSB) register of the generator.

The 9-bits start value of the scrambling generator should be 101010101.

3.4.1 Rationale

Scrambling is very helpful to avoid long sequences of only zeros or only ones as these could make bit synchronization unnecessarily complicated.

[1] NRZI (i.e. differential encoding) as specified in the current VDL4 standard is not required for correct signal detection and demodulation.

3.5 Burst Formats and Error Control

The following suggestion defines all three burst formats used in VDL4 (single slot, half slot, and multi-slot)[2].

In all figures the length unit is *byte*.

3.5.1 Single-Slot Burst

The figure below shows the new format of the single slot burst.

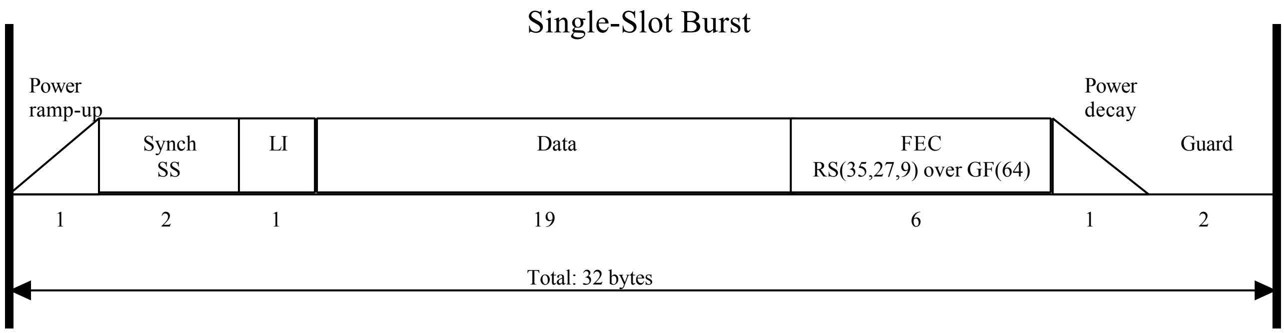

The synchonization sequence is followed by the LI (8 bits) and the data segment. In single slots the LI can take on the values [1, 19] and indicates the number of used data bytes in the slot.

A maximum rate of 11.4 kbps can be obtained when LI =19.

The data segment consists of LI bytes which can be used entirely for the payload (information) as there is no bit stuffing. Unused data bytes are set to zero (padding).

The forward error correction (FEC) segment consists of 6 bytes. The code is a Reed-Solomon (RS) code over the Galois field GF(64). The code works on 6-bits symbols and is defined by its blocklength N, dimension K and minimum distance D. The RS code protects the LI and the data.

The code has parameters N = 35 symbols, K = 27 symbols and D = 9 and can correct up to 4 symbol errors.

As (1 + 19)*8 = 6 x 26 + 4 is not a multiple of 6 the last two (2) bits of the last information symbol are not used; these two bits are assumed to be zero and shall not be transmitted.

The RS code is used for combined error correction and detection. The RS code's performance is described in detail in section 3.5.5.

3.5.2 Half-Slot Burst (Delayed Burst)

The figure below shows the new format of the half-slot burst.

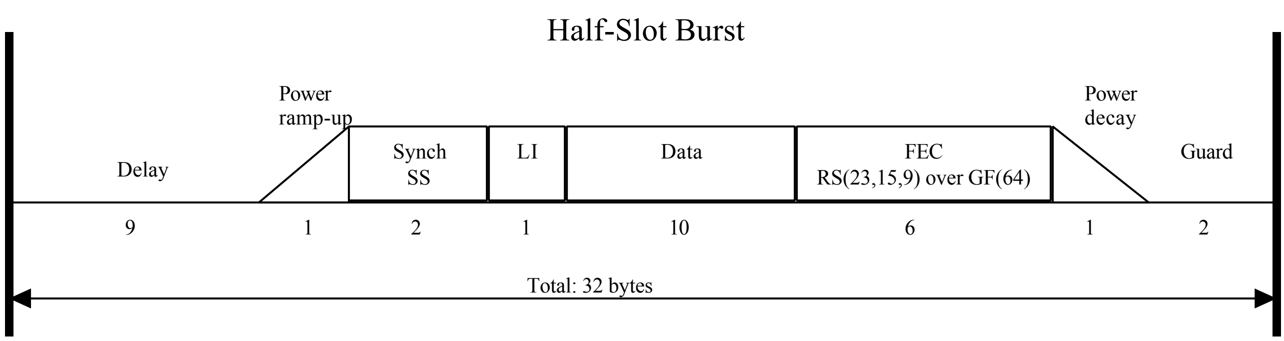

The slot starts after 9 bytes of delay which corresponds to 3750 microseconds (which corresponds to ~ 606 nmi).

---

[2] HDLC is removed in this suggestion and is thus not included in the suggested burst formats.

The synchonization sequence is followed by the LI (8 bits) and the data segment. In half-slot bursts the LI can take on the values [1, 10].

The data segment consists of LI bytes which can be used entirely for the payload (information) as there is no bit stuffing. Unused data bytes are set to zero (padding).

The forward error correction segment consists of 6 bytes. The code is a Reed-Solomon code over the Galois field GF(64). The code works on 6-bits symbols and is defined by its blocklength N, dimension K and minimum distance D. The RS code protects the LI and the data.

The code has parameters N = 23 symbols, K = 15 symbols and D = 9 and can correct up to 4 symbol errors.

As (1 + 10)*8 = 6 x 14 + 4 is not a multiple of 6 the last two (2) bits of the last information symbol are not used; these two bits are assumed to be zero and shall not be transmitted.

The RS code is used for combined error correction and detection. The RS code's performance is described in detail in section 3.5.5.

3.5.3 Multi-Slot Burst

The figure below shows the new format of the multi-slot burst.

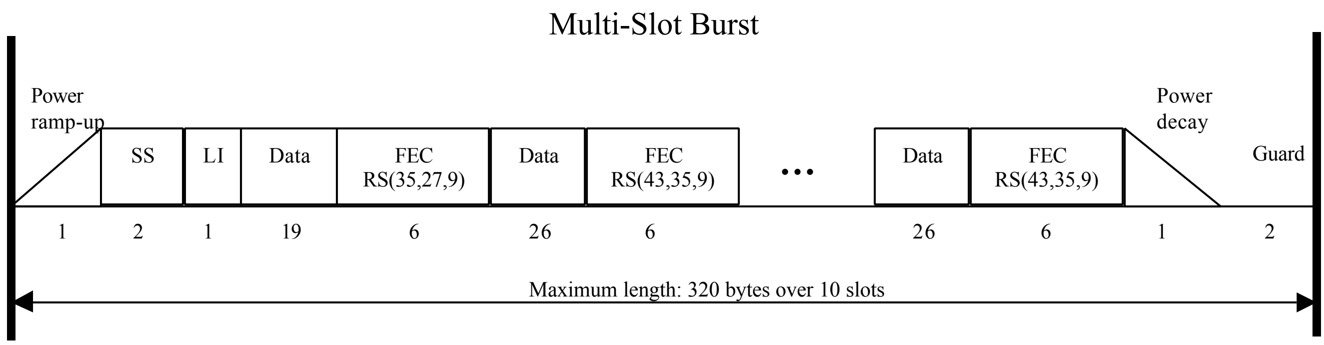

The LI indicates the total number of data bytes (excluding the FEC) in the whole multi-slot burst.

In multi-slot bursts the LI can take on the values [20, 253] (see table in 3.5.4).

A maximum of 10 slots can be used which implies a maximum message length of 19 + 9x26 = 253 bytes equivalent to a rate of 15.18 kbps. The shortest multi-slot burst consists of two slots.

The data consist of one 19-bytes segment plus 1-9 segments of at most 26 bytes which can be used entirely for the payload (information) as there is no bit stuffing. Unused data bytes in the last segment are set to zero (padding).

A multi-slot burst occupies 2-10 slots.

The number of slots occupied is denoted by S and is obtained from the following formula:

$S = 1 + \lceil (LI - 19) / 26 \rceil$

where $\lceil x \rceil$ is the roof of x.

The FEC segments consists of 6 bytes in all codewords. The code is a Reed-Solomon (RS) code over the Galois field GF(64). The code works on 6-bits symbols and is defined by its blocklength N, dimension and minimum distance D. In the first codeword the RS code protects the LI and the data.

The codes have parameters N = 35, 43 symbols, K = 27, 35 symbols and D = 9 and can correct up to 4 symbol errors. It is in fact the same RS(63,55,9) code shortened in different ways (it is also the same RS code as in the single and half-slot bursts).

As (1+19)*8 = 160 and 26*8 = 208 are not multiples of 6 the last two (2) bits of the last information symbol are not used; these two bits are assumed to be zero and shall not be transmitted.

The RS code is used for combined error correction and detection. The RS code's performance is described in detail in section 3.5.5.

3.5.4 Coding of the Length Indicator LI

The LI is coded as shown in the table below.

LI Value	Interpretation	Comment
0	RFU	Reserved
1-19	Number of data bytes in a Single or Half slot burst	
20-253	Number of data bytes in a Multi-slot burst	
254-255	RFU	Reserved

3.5.5 Performance of the RS Codes

RS codes are excellent for combined error correction and detection.

The *guaranteed* performance of the suggested RS code is shown in the table below. In the table the code is used for correction up to one symbol less than the maximum error correction capability in order to allow for a proper amount of error detection.

The RS(23,15,9), RS(35,27,9) and RS(43,35,9) codes are all shortened versions of the same RS(63,55,9) code over GF(64) and can thus be encoded/decoded by the same software or hardware implementation.

The code is constructed over the Galois field GF(64) as generated by the primitive polynomial $p(x) = 1 + x + x^6$. An RS symbol is thus 6 bits long.

The code's generator polynomial is as follows:

$$G(x) = \prod_{i=m}^{m+7} (x - a^{i})$$

where α is a (primitive) root of p(x) and $\mu = 28$ (this value of μ makes G(x) self-reciprocal).

Code	Maximum random error correction	Applied random error correction	Applied random error detection	Applied burst error correction	Applied burst error detection	Comment
RS(23,15,9) RS(35,27,9) RS(43,35,9)	4 symbols	1-3 symbols	4 –5 symbols	1-3 symbols e.g. all bursts of 1-13 bits	4-5 symbols e.g. all bursts of 14-25 bits	Many bursts of length 14-15 bits will also be corrected. Many bursts of length 26-27 bits will also be detected.

Table 1. Performance of the suggested RS code.

The suggested code allows for flexible use of error correction and detection while being easily implemented in hardware or on DSPs.

3.5.6 Rationale

- Forward error control codes (FEC) are required due to the non-ideal characteristics of the radio channel. In fading channels such as the VDL4 radio channel the BER will stay rather high also at very large SNR values.
- FEC improves the CCI behavior of the radio channel.
- The multi-slot format gives no degradation in BER performance compared to single slot

3.6 Guard Time

The guard time should be 16 bits (~ 832 microseconds) long.

3.6.1 Rationale

This guard time plus the transmitter power decay time add to 24 bits or ~1250 microseconds which is equivalent to ~ 202 nmi.

In the worst case two units 200 nmi apart will interact as shown in figure below.

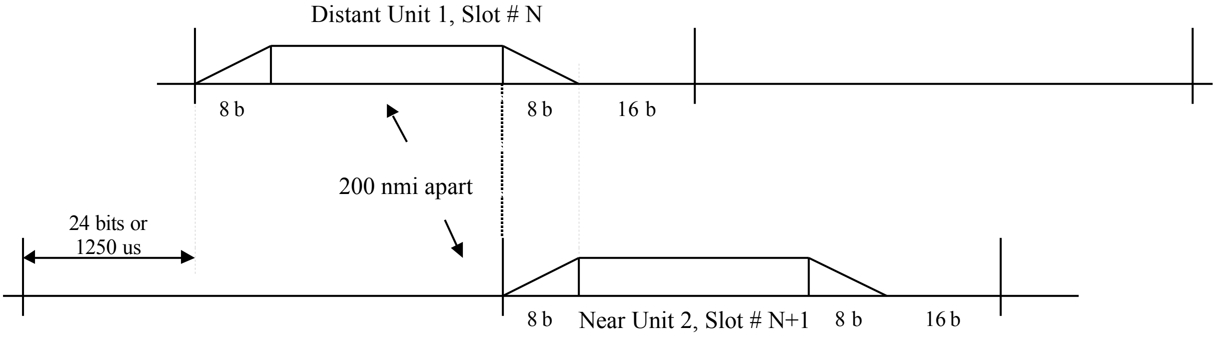

Notice that no data is transmitted during power ramp-up/decay.
Also the large distance between the units implies a very small received power level during the second half of the power decay period (which overlaps with the second half of the power ramp up period of the next slot).

Thanks to guard space and FEC the performance in failure mode will be significantly improved.