D3.2 SpaceWire-RT Updated Specification

Lead Beneficiary:	University of Dundee		
Author(s):	Steve Parkes, UNIVDUN		
Work Package:	WP3	**Task:**	Task 6
Version:	1.00	**Last modified:**	6/2/2013
Status:	Issued		
Approved by:	S.M. Parkes	**Date:**	6th February 2013

Disclaimer:

The material contained in this document is provided for information purposes only. No warranty is given in relation to use that may be made of it and neither the copyright owners or the European Commission accept any liability for loss or damage to a third party arising from such use.

Copyright Notice:

Table of Contents

Table of Figures

Table of Tables

1 Introduction

This document records changes made to the SpaceFibre standard following the extensive simulation work carried out in Work Package 3 and reported in WD 3.1 “SpaceWire-RT SDL and System-C Simulation Models” and WD 3.2 “SpaceWire-RT SDL and System-C Simulation Results”.

1.1 SpaceWire-RT Objectives Addressed

This report addresses the following objectives of the SpaceWire-RT project:

- A validated SpaceWire-RT specification with important, novel features of SpaceWire-RT networks tested using SDL models.
- A draft standard document for SpaceWire-RT, which has been reviewed by the International SpaceWire Working Group.
- Disseminated results of the SpaceWire-RT study to the European and Russian space industries, and to the international space community.

In particular this report provides an update to the SpaceFibre standard, which includes the novel features provided as a result of the SpaceWire-RT research and which takes into account feedback from SUAI on the standard following their simulation activities.

In addition this report outlines further analysis of the SpaceWire over SpaceFibre concept, which has resulted in this approach being removed from the SpaceWire-RT protocol stack

1.2 Guide to the Report

Section 2 describes the layers that have been added to the SpaceFibre standard as a result of the SpaceWire-RT project.

Section 3 examines the SpaceFibre over SpaceWire concept, concluding that it is not possible to provide the SpaceFibre FDIR capabilities over a SpaceWire network without designing specific routing switch devices.

Section 4 presents an update to the SpaceFibre protocol stack, adding network and management layers, and removing the SpaceFibre over SpaceWire protocols.

Section 5 is the main part of this report which presents the change requests made by SUAI following the simulation work in WP3. These change requests are analysed and comments made and ensuing actions determined.

Section 6 is an update to the SpaceWire-RT outline specification initially presented in D2.1.

Section 7 present the conclusions of this report and outlines the next/concurrent stages of the SpaceWire-RT project.

Annex 1 provides the SpaceFibre standard draft F 0.1, which has been updated following the change request actions described in section 5.

1.3 References

[1] Parkes SM, Ferrer Florit A, Gonzalez A, and McClements C, “SpaceFibre”, Draft D, Space Technology Centre, University of Dundee, 29th February 2012.

[2] Parkes SM, Ferrer Florit A, Gonzalez A, and McClements C, “SpaceFibre”, Draft E1, Space Technology Centre, University of Dundee, 28th September 2012.

[3] Yuriy Sheynin, Elena Suvorova, Valentin Olenev, Irina Lavrovskaya, Ilya Korobkov, Pavel Morozkin, and Nikolay Sinyov, “SpaceWire-RT SDL and System-C Simulation Models”, SpaceWire-RT project report WD3.1, SUAI, 18th January 2013.

[4] Yuriy Sheynin, Elena Suvorova, Valentin Olenev, Irina Lavrovskaya, Ilya Korobkov, Pavel Morozkin, and Nikolay Sinyov, “SpaceWire-RT SDL and System-C Simulation Results”, SpaceWire-RT project report WD3.2, SUAI, February 2013.

[5] A. X. Widmer and P. A. Franaszek, "A DC-Balanced, partitioned-Block, 8B/10B Transmission Code", IBM Journal of Research and Development archive, Volume 27 Issue 5, September 1983, Pages 440-451.

2 Layers Added to the SpaceFibre Standard

The SpaceWire-RT project builds on the emerging SpaceFibre standard which is being designed by the University of Dundee with inputs from international spacecraft engineers. The aim is to publish the final standard though the European Cooperation for Space Standardization (ECSS). The SpaceWire-RT project has contributed to the following parts of the SpaceFibre standard:

- QoS Mechanisms in the Virtual Channel layer,
- FDIR Mechanisms in the Retry layer,
- Network level concepts in the Network layer.

A description of the research in these three areas is available in sections 6, 7, and 12 respectively of the "D2.1 SpaceWire-RT Outline Specification" SpaceWire-RT report.

The QoS and FDIR work has been fully integrated with the SpaceFibre draft standard (Draft E1) resulting in sections 5.4.4 "Media Access Control" and 5.7 "Retry Layer". The work on the Network layer has yet to be incorporated in the draft standard.

The simulation carried out in Work Package 3 of the SpaceWire-RT project covered the following versions of the SpaceFibre draft standard:

- SDL Modelling
    - Encoding and Lane Layers – Draft D
    - Retry, Frame, Virtual Channel and Broadcast Channel Layers – Draft E1
- System-C Modelling
    - Draft E1

The latest version of the standard (draft E1) was used for the critical simulation of the Virtual Channel and Retry Layers. Draft D was used for the lower layers because this simulation work was carried out prior to Draft E1 of the SpaceFibre standard being written.

3 SpaceFibre Over SpaceWire

In this section the SpaceFibre over SpaceWire concept is analysed.

3.1 SpaceFibre Over SpaceWire

In D2.1 the possibility of running SpaceFibre higher layers over a SpaceWire lane layer was considered. The lane layer of SpaceFibre would be replaced by SpaceWire, with SpaceFibre frames being embedded in SpaceWire packets. The corresponding protocol stack is illustrated in Figure 3-1.

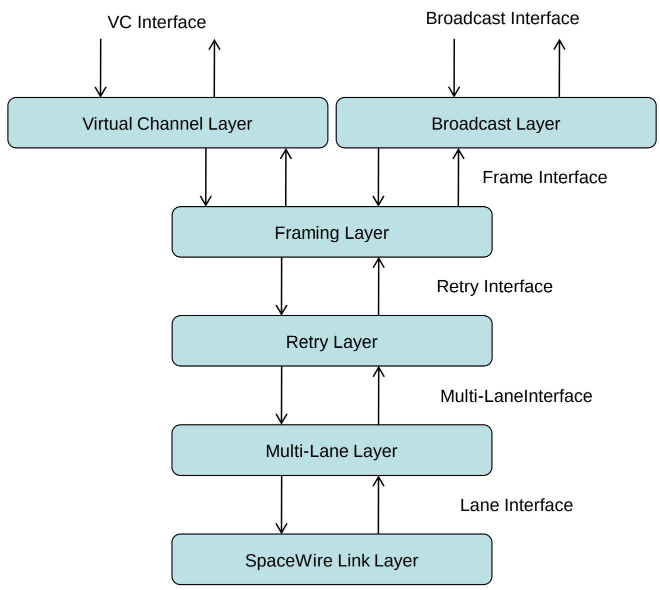

Figure 3-1 SpaceFibre Over SpaceWire

The SpaceFibre lane, encoding, serialisation and physical layers are replaced by the SpaceWire Link Layer, which is standard SpaceWire. The user application passes SpaceWire packets to be sent over the network into the virtual channel layer. These packets are segmented into frames by the framing layer. The frames are passed to the SpaceWire link level where they are encapsulated into individual SpaceWire packets. Since the maximum size of a SpaceWire packet is 66 words (264 bytes) the SpaceWire packets are all small. Multiple application packets can be multiplexed over a single SpaceWire packet using the SpaceFibre virtual channels and medium access controller.

At the link level this would work fine, but at the network level it requires new SpaceFibre/SpaceWire routers to be designed. If this is to be done there is little benefit in the SpaceFibre over SpaceWire approach.

One other possibility is to run SpaceFibre higher level protocols over a complete SpaceWire network, as illustrated in Figure 3-2.

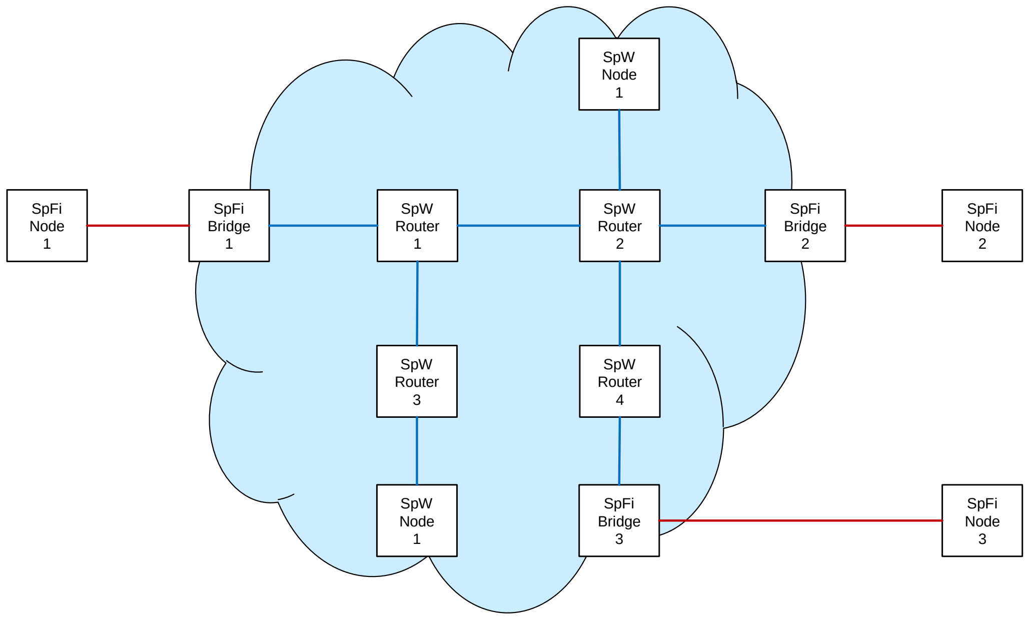

Figure 3-2 SpaceFibre Over SpaceWire Network

The intention is that a SpaceFibre Bridge device (e.g. SpFi Bridge 1) would pack SpaceFibre frames and control codes into SpaceWire packets and send them over the SpaceWire network (highlighted by blue cloud) to another SpaceFibre Bridge device (e.g. SpFi Bridge 2) where the SpaceFibre Bridge would unpack them reconstituting the SpaceFibre traffic. This would then allow SpaceFibre higher level protocols like QoS to run over SpaceWire. Furthermore different virtual channels could be routed to different SpaceFibre bridge devices, making use of SpaceWire routing. This would then allow one SpFi node to talk to VCs in other SpFi nodes attached via SpFi bridge devices to the SpaceWire network.

For QoS and flow-control, this would work fine, as there is a separate buffering for each virtual channel. Flow control, for example, would then operate between the two SpaceFibre nodes attached via the SpaceWire network. The data frames and FCTs related to a particular virtual channel would be given appropriate SpaceWire path or logical addresses to route them to the correct SpaceFibre Bridge and then on to the destination SpaceFibre node.

However, the retry mechanism in SpaceFibre has been designed to use a single retry buffer for all virtual channels, in order to minimise the overall amount of buffer space required. So retry control words operate over a link only and do not know anything about the virtual channel. ACKs and NACKs are per link, not per virtual channel. This means that it is not possible to run the SpaceFibre retry protocols over the SpaceWire network. It is possible over a single SpaceWire link, but this is not really of very much use.

Because it is not possible to operate the SpaceFibre retry protocols bridged over a SpaceWire network, the SpaceFibre over SpaceWire concept is not considered further. Connection between SpaceWire and SpaceFibre network can be achieved with a simple bridge device

3.2 SpaceWire to SpaceFibre Bridge

The aim of a SpaceWire to SpaceFibre Bridge is to be able to connect a legacy SpaceWire device to a SpaceFibre network. This is simple to do with SpaceFibre: each SpaceFibre virtual channel can be connected to a SpaceWire link as illustrated in Figure 3-3.

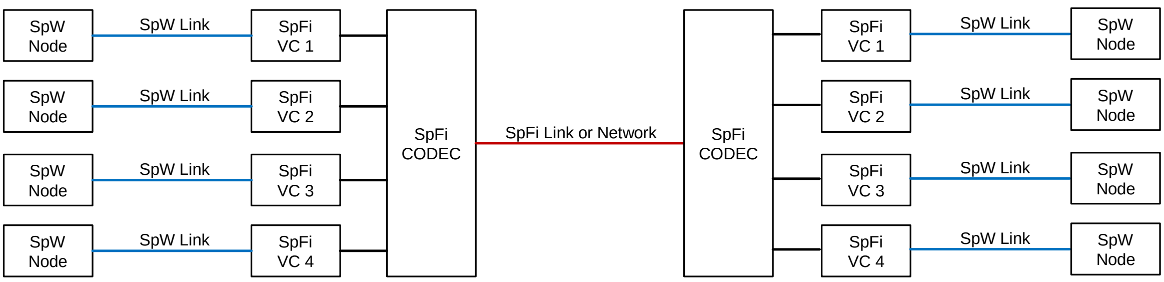

Figure 3-3 SpaceWire to SpaceFibre Bridge

SpaceFibre virtual channels send and receive SpaceWire packets, making it simple to bridge SpaceWire over SpaceFibre. A SpaceWire link can be connected to a SpaceFibre virtual channel. In Figure 3-3 four SpaceWire nodes are connected to virtual channels in a SpaceFibre CODEC. The SpaceFibre CODEC is attached via a SpaceFibre link to another SpaceFibre CODEC. SpaceWire packets from a SpaceWire node travel over a SpaceWire link to the SpaceFibre virtual channel. The data from each virtual channel is multiplexed over the SpaceFibre link according to the quality of service settings for each virtual channel. The SpaceWire packets emerge at the corresponding virtual channel at the far end of the SpaceFibre link. The received SpaceWire packets are then send over the SpaceWire link to the destination SpaceWire node.

SpaceWire links operate at up to 200 Mbits/s (typical maximum for spaceflight implementation). This means that up to 12 SpaceWire links can be multiplexed over a single SpaceFibre link. Using a SpaceWire to SpaceFibre Bridge to concentrate traffic over a single SpaceFibre link can save significant cable mass on-board a spacecraft.

Note that SpaceWire time-codes have to be handled separately to the SpaceWire packet data and can be sent over SpaceFibre using broadcast messages.

4 Updated SpaceFibre Protocol Stack

The SpaceWire-RT protocols are illustrated in the combined protocol stack of Figure 4-1.

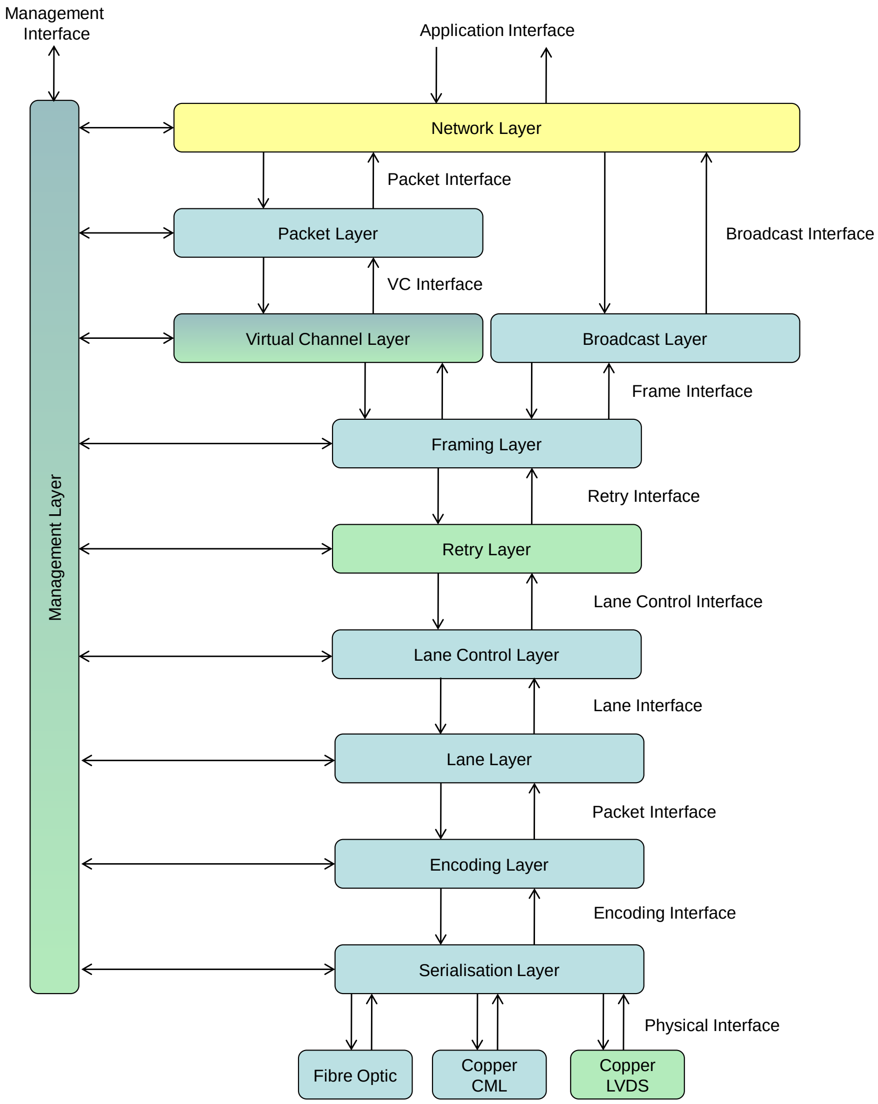

Figure 4-1 SpaceWire-RT Protocol Stack

The blue elements in Figure 4-1 are the layers initially defined by SpaceFibre. The green elements are those that have been designed within the SpaceWire-RT project. The Network Management and Virtual Channel layer have been partly covered by the SpaceWire-RT project. The Network layer is not yet included in the SpaceFibre standard.

There are twelve conceptual layers to the SpaceFibre:

- Network: responsible for routing SpaceFibre packets over a SpaceFibre network, comprising SpaceFibre routing switches, SpaceFibre links, and SpaceFibre nodes. Also responsible for validating and broadcasting SpaceFibre broadcast messages over a SpaceFibre network.
- Packet: responsible for forming data to be sent over a SpaceFibre link into packets, which have the same format as SpaceWire packets.
- Virtual Channel and Flow Control: responsible for quality of service and flow control over the SpaceFibre link.
- Broadcast: responsible for sending short broadcast messages across a SpaceFibre link and for receiving those messages.
- Framing: responsible for framing SpaceFibre packet data, broadcast messages and FCTs to be sent over the SpaceFibre link. It is also responsible for scrambling SpaceWire packet data for EMC mitigation purposes.
- Retry: responsible for recovering from transient errors on the SpaceFibre link, and for reporting errors and link failure. Detects missing and out of sequence frames.
- Multi-Lane: responsible for operating several SpaceFibre lanes in parallel to provide a higher data throughput and to provide redundancy with graceful degradation.
- Lane: responsible for initialising the lane, detecting lane errors and re-initialising the lane after an error has been detected.
- Encoding: responsible for encoding the data and control words into a suitable form for sending over the SpaceFibre link and decoding received data and control words. Uses 8B/10B encoding.
- Serialisation: responsible for serialising and de-serialising encoded data and control words for sending and receiving over the serial interfaces.
- Physical: responsible for sending the SpaceFibre information over the physical media used in SpaceFibre: fibre optic, Current Mode Logic (CML) and Low Voltage Differential Signalling (LVDS). Fibre Channel copper medium may also be added in future.
- Management: responsible for configuring, controlling and monitoring the status of the various layers of the SpaceFibre protocol stack. This can be done by a local or remote network management application.

The key characteristics of the different SpaceWire-RT protocols, which are differentiated by their physical layer, are detailed in Table 4-1.

Table 4-1 SpaceWire-RT Protocol Characteristics

Characteristic	**SpFi-FO**	**SpFi-CML**	**SpFi- LVDS**
Media	Fibre Optic	Copper CML	Copper LVDS
Encoding	8B/10B	8B/10B	8B/10B
Speed Range	0.1 to 20 Gbits/s 50 Gbits/s in future	0.1 to 20 Gbits/s 50 Gbits/s in future	1 to 600 Mbits/s 1 to 100 Mbits/s OS
Distance	100 m	5 m	10 m
Galvanic Isolation	Yes	Yes	Yes
Packet Size	Arbitrary	Arbitrary	Arbitrary
SpaceWire Packet Level	Yes	Yes	Yes
Latency (TBC)	<0.5 µs	<0.5 µs	1 µs
Cable Mass	< 30g/m	< 30g/m	< 30g/m
Power (TBC)	< 200 mW	< 200 mW	< 200 mW
QoS BW Reserved	Yes	Yes	Yes
QoS Priority	Yes	Yes	Yes
QoS Scheduled	Yes	Yes	Yes
QoS Best Effort	Yes	Yes	Yes
Broadcast Message	Yes	Yes	Yes
Determinism	Yes	Yes	Yes
Reliability	Yes	Yes	Yes
Fault Detection	Yes	Yes	Yes
Fault Isolation	Yes	Yes	Yes
Retry	Yes	Yes	Yes
SpaceWire compatible	No	No	No

5 Requested Changes to SpaceFibre Standard

In this section the changes to the SpaceFibre standard draft E1 requested by SUAI as an output from WP3 are recorded and analysed. Each level of the standard is covered in the following subsections. The table provides the changes requested referencing the specific clause of the SpaceFibre standard draft E1. The reply to the change request and any action to be taken is added to each change request. The modifications to the standard provide in the Action/Comment column of the following tables, are provided in draft F 0.1 of the SpaceFibre standard which is included as Annex 1 of this report.

5.1 Formats

#	Page	Section/Figure/Table #	Change Request	Action/Comment
1	78	Section 5.3.3.2 f.	SIF does not contain the Seed field. Seed is another word which is a part of the idle-frame.	The Seed Field is the first word after the SIF. This has been corrected. A CRC has been added to the SIF to protect its frame sequence number.
2	79	Section 5.3.3.2 c. (NOTE)	FCT can also terminate an idle-frame. Also (N)ACK terminates IDLE frame.	Idle frames shall be treated in the same way as data frames: FCTs, ACK, NACKs, etc. are embedded in them rather than terminating the frame. This means that idle frames replace the flow of data frames when there is no data to transfer and FCTs, ACKs, NACKs are handled in the same way regardless of whether it is a data frame or idle frame that is being transmitted. The standard has been updated accordingly.
3	-	The whole specification	The SpaceFibre uses two different commas for different purposes. Init comma K28.5 – for initialisation of a link, and comma K28.7 – for all other purposes. However, the need of two comma codes is not obvious, while it	There is a good reason for using two different commas. The initialisation comma K28.5 has disparity and a valid inverse This enables it to be used during

			is possible to use only one comma. Moreover, the K-code K28.7, followed by certain other data or control codes can produce a false comma, which can cause some errors in functionality, which can be avoided by the use of K28.5 comma only.	initialisation when the input may be inverted allowing the clock-data recovery circuit to lock regardless of the input polarity. The K28.7 has even disparity which when sending INIT words means it is never changing the disparity. That is why it is not being used for initialization. However, the K28.7 has the important property that it cannot be generated from a data word or any other used symbol when a single bit flip error occurs. This ensures that we can prove that the protocol is robust under any single bit flip error. The symbols selected to be used after the K28.7 have been carefully selected, so a false comma is never produced. No change to the standard. Reference: A. X. Widmer and P. A. Franaszek, "A DC-Balanced, partitioned-Block, 8B/10B Transmission Code", IBM Journal of Research and Development archive, Volume 27 Issue 5, September 1983, Pages 440-451.

5.2 Virtual Channel layer

#	Page	Section/Figure/Table #	CR	Type
1	84	Section 5.4.1 h.	It defines that *"it shall be possible to start writing the next SpaceWire packet to be sent over that virtual channel, provided that there is room in the output virtual channel buffer for at least one more N-Char"*. This	SpaceWire can also be considered a byte stream interface. SpaceFibre is compatible with SpaceWire. A packet layer description will be added to the SpaceFibre standard to

#	Page	Section/Figure/Table #	CR	Type
			point results in byte stream interface, not in packet interface, which is defined for the SpaceFibre CODEC.	clarify this conceptual concern. This change has not yet been made to the standard.
2	84	Section 5.4.1 p.	According to this point, all output virtual channel buffers should be flushed. However, it seems to be unnecessary as it can lead to deletion of several entire SpW packets. So we propose the following solution: for each virtual channel in the output buffer all the data should be removed starting from the beginning of the buffer FIFO including the first EOP. All other data characters should be left in the buffer. If the output buffer became empty after deletion and if the last character written by the host application was not an EOP, all new data characters will be discarded up to and including the next EOP.	Remote flush is a serious condition. It should rarely happen and indicates a major concern. Considering this, and for the sake of simplicity, it is simpler and safer to empty the buffers and discard input data until the next EOP marker. No change to the standard.
3	85	Section 5.4.1 q.	According to this point, all input virtual channel buffers should be flushed. However, it seems to be unnecessary as it can lead to deletion of several entire SpW packets, which were received correctly. So we propose the following solution: for each virtual channel in the input buffer all the data should be removed starting from the end of the buffer FIFO including the first EOP. All other data characters should be left in the buffer. If the input buffer became empty after deletion and if the last character which is read by the host application was not an EOP, an EEP shall be the next symbol read by the host application.	As the previous comments, it is safer and simpler to spill everything and insert EEP if required. Data will probably be lost almost certainly anyway. For example, an unexpected cold reset will produce a massive loss of data, so a few more packets lost makes no difference. No change to standard.
4	88	Section 5.4.4.2 f (NOTE)	There is no "quality of service"	It is the priority level quality of

#	Page	Section/Figure/Table #	CR	Type
			management configuration parameter (see Table 5-43)	service parameter. The standard has been updated to clarify this.
5	89	Table 5-11	According to the table the virtual channel with best effort QoS has its priority quality of service parameter set to the lowest priority. However, the default value for this parameter is 0 (see Table 5-43), which corresponds to the highest priority level. Therefore, the default value for the configuration parameter should be changed to 15.	Yes, this is a mistake in the standard. The default priority level should be 15 resulting in the best effort quality of service. The standard has been updated.
6	90	Section 5.4.4.3 b.	It is not clear, how the Bandwidth Credit can be updated if no data frame was send during the interval approximately equal to the time it takes to the send a full data frame.	This condition is set to define the minimum update rate of the bandwidth counters when nothing must be sent. The clause means that every 66 words or less the different counters should be updated. A note has been added to the standard to clarify this.
7	90	Section 5.4.4.3 c.	As far as we understand, the Available Bandwidth cannot take into account the overheads due to framing and control words.	It is impossible to know exactly the bandwidth used by the overheads but a worst case estimation can be made. A margin should be allocated for broadcast messages and each virtual channel should be allocated a little more bandwidth than it is actually expect to use. This will result in some potentially unused bandwidth, but in practice this can be taken up by any best effort virtual channels, or other virtual channels that have the possibility of sending more data than specified by their bandwidth allocation.

#	Page	Section/Figure/Table #	CR	Type
				When specifying a bandwidth usage for a virtual channel it is simpler for the user to think in terms of a fraction of link capacity. Otherwise it is difficult to take into account the bandwidth used by the broadcast. The total allocated bandwidth for all the virtual channels must be less than 100% to allow for overheads and broadcast messages. No change to the standard.
8	92	Section 5.4.4.3 k. (NOTE)	It should be specified how it can be prevented.	By “prevent a virtual channel sending more data” it is meant “stop a virtual channel sending more data”, i.e. the virtual channel is disabled stopping it from transmitting any more data. The standard has been updated clarifying this note.
9	92	Section 5.4.4.3 n.	1. This leads to the fact that upon a cold reset no virtual channel is able to send any data. However, it can be necessary for network configuration. The following solution can be more reasonable: *"Upon cold reset the Expected Bandwidth for each virtual channel shall be set to 1/N, where N is the number of virtual channels."* 2. Moreover, the simulation shown, that if one VC of all is not assigned any value for expected bandwidth, it will lead to division by zero, while updating Bandwidth Credit Counter.	The point about division by zero is well made. It shall not be possible to set the Expected Percentage Bandwidth to zero; there will be a minimum value to which it can be set. A separate control shall be used to fully disable a virtual channel. On cold reset the Expected Percentage Bandwidth for each virtual channel shall be set to the minimum value, except for VC0 which shall have its Expected Percentage Bandwidth set to 10%. Note that VC0 is used for network configuration, control and monitoring. The standard has been updated accordingly.

#	Page	Section/Figure/Table #	CR	Type
10	159	Table 5-43	The default value assigned to virtual channels is defined as '0'. However, this corresponds to the highest priority level and cannot be used as default value.	See reply to CR #5.
11	159	Table 5-43	The quality of service parameter shall be included in this table.	The quality of service is derived from the values inserted in the Priority level, Expected Bandwidth and Allocated Time Slots parameters. This has been clarified in the standard.
12	159	Table 5-43	The title of the table should be changed from *"SpaceWire CODEC Configuration Parameters"* to *"SpaceFibre CODEC Configuration Parameters".*	The standard has been updated as requested.

5.3 Broadcast Layer

#	Page	Section/Figure/Table #	CR	Type
1	56	Section 5.2.2.3.2 (NOTE 2)	What numerical values correspond to the *"All"* broadcast channels and *"All"* broadcast message types while registering to receive broadcast messages?	These parts of the standard belong to the network layer and have been removed from the current version of the standard to avoid confusion.
2	95	Section 5.5.6 a.	Broadcast channel 7 is reserved for synchronisation broadcast. Network management broadcast channels are 8 to 31. Therefore, 7 should be replaced by 8 in this clause	These parts of the standard belong to the network layer and have been removed from the current version of the standard to avoid confusion.

5.4 Framing Layer

#	Page	Section/Figure/Table #	CR	Type
1	99	Section 5.6.3.1 a.	This point defines bit-wise multiplication to be used while data scrambling. However, data scrambling is performed by XORing with a sequence of random numbers produced from a scrambling polynomial. It will be better to add this information to this point.	The standard has been updated as requested.
2	99	Section 5.6.3.1 f.	What should be done with the random number while the EOP, EEP or Fill occurred in the data stream? Should it wait for another data symbol without any actions?	The random number generator should continue with its normal operation even when its result is not needed because the character is an EOP, EEP or Fill. The standard has been updated accordingly.
3	31	Section 4.3.4	In the third paragraph it is stated that the framing layer passes data frames, FCTs and broadcast frames to the retry layer in the order in which they are to be sent over the SpaceFibre link. However, this contradicts the specification of the framing layer, which is not responsible for this task. It is the retry layer to arbitrate data flows.	The standard has been clarified.

5.5 Retry Layer

#	Page	Section/Figure/Table #	CR	Type
1	80	Section 5.3.4 c.	The difference in precedence of ACK and NACK should be defined.	If a NACK is to be send no ACK will be sent and vice versa

#	Page	Section/Figure/Table #	CR	Type
				This has been clarified in the standard
2	81	Section 5.3.4 o.	FULL cannot be inserted within a data frame, as it has lower precedence than a data frame in accordance with point c.	This is correct FULL has lower precedence than a data frame and cannot be inserted within a data frame. The standard has been corrected.
3	100	Section 5.7.1.2.1	The statement *"The retry layer shall pass a TX_DATA_RDY.indication primitive to the framing layer to indicate that it is ready to accept a new data or control word"* should be changed to *"The retry layer shall pass a TX_DATA_RDY.indication primitive to the framing layer to indicate that it is ready to accept a <u>new data frame</u>"*	The service interface specifications in the standard will be reviewed and checked for consistency.
4	101	Section 5.7.1.2.3	The statement *"The TX_DATA_RDY.indication primitive shall be passed to the framing layer when the lane layer can accept a new request to transfer a data frame."* should be changed to *"The TX_DATA_RDY.indication primitive shall be passed to the framing layer when the <u>retry</u> layer can accept a new request to transfer a data frame."*	The standard has been updated as requested.
5	102	Section 5.7.1.5.3	The statement *"The TX_FCT_RDY.indication primitive shall be passed to the framing layer when the lane layer can accept a new request to transfer an FCT"* should be changed to "*The TX_FCT_RDY.indication primitive shall be passed to the framing layer when the <u>retry</u> layer can accept a new request to transfer an FCT."*	The standard has been updated as requested.
6	103	Section 5.7.1.8.3	The statement *"The TX_BC_RDY.indication primitive shall be passed to the framing layer when the lane layer can accept a new request to transfer a broadcast frame.*" should be	The standard has been updated as requested.

#	Page	Section/Figure/Table #	CR	Type
			changed to "*The TX_BC_RDY.indication primitive shall be passed to the framing layer when the retry layer can accept a new request to transfer a broadcast frame.*"	
7	104	Figure 5-7	To the exit conditions of the states RxDataFrame, RxBroadcastFrame, RxBroadcast&DataFrame and RxIdleFrame in case of Frame Error, RXERR should be added.	Clarifications have been made to the Data Word Identification State Machine (Figure 5-7).
8	105	Section 5.7.2.1 c.	How should be processed the reception of EDF and EBF?	No action must be done in this case. A note has been added for clarification.
9	106, 107	Sections 5.7.2.3 c., 5.7.2.4 c.	How should be processed the reception of SBF control word in RxBroadcastFrame state?	It's an error. It has been described in the figure but not in the text. The standard has been updated.
10	109	Section 5.7.2.6	1. The RETRY control word has higher precedence than broadcast frame. Consequently, it can be received in RxBroadcastFrame and RxBroadcast&DataFrame states. 2. Consider the case that FCT, ACK, NACK or FULL are received in RxBroadcastFrame or RxBroadcast&DataFrame states and they are valid. What should be done then? Should they be rejected (it is especially important for ACKs and NACKs)? Should the Data Word Synchronisation state machine move to the RxNothing state and flush the buffers?	If there are no errors the sender should not add control words within a broadcast frame. In the receiver the control words (FCT,ACK,NACK,FULL) are always processed in the same way regardless of the state of the word identification machine. Reception of a RETRY control word now affects the operation of the data word identification state machine. These modifications have been included in the SpaceFibre standard.
11	111	Section 5.7.5 e.1.(c)	There is a mistake in polynomial coefficients in the second component. It should be *n-1*	It should be n-2. The standard has been updated accordingly.

#	Page	Section/Figure/Table #	CR	Type
12	114	Section 5.7.6.2 l. and NOTE	The received 7-bit frame sequence count field of a received FULL shall not be the same value as the receive frame sequence counter plus one. It shall be equal to the last received frame sequence counter. Consequently, FULL should be removed from this point.	The standard has been updated as requested.
13	114	Section 5.7.6.2 m.	FULL shall be added to this point.	The standard has been updated as requested.
14	116	Section 5.7.7 k.	According to this point the number of retries status parameter should be added to the Table 5-44	The standard has been updated to include this parameter in the table.
15	118	Section 5.7.9.2 d.	When a NACK is sent it shall contain not inverted receive polarity flag. Otherwise, if it contains the receive polarity flag inverted, it will never be considered as valid. See the diagram, with the explanation. 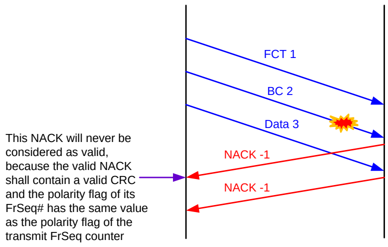 	Internally the sequence counter has already been inverted. That's why we have to send the value inverted, as this will originate the NACK + 1 to be sent (not NACK -1). Clarifications have been made to the standard.
16	118	Section 5.7.9.2 c.	How NACK can be pending? According to the section 5.7.9.2 there is no rule which will cause pending of NACK control word.	The sentence "A NACK shall be requested to be sent" means that a NACK will be pending to be sent if it cannot be sent immediately. For example, if FULL words arrive all the time and they have the wrong sequence number, theoretically a NACK should be sent for any of those FULL words. Hence, continuous sending of NACK words would appear in the other end. To prevent this situation there is a window for ACKs and NACKs, and they can be sent only once every 16

#	Page	Section/Figure/Table #	CR	Type
				words (if there is a change between ACK and NACK then it will be sent immediately). The standard has been clarified accordingly.

5.6 Lane Layer

#	Page	Section/Figure/Table #	CR	Type
1	129	Section 5.9.2.2 a.7.	Should these 8 LOS words be consecutive? They are defined as consecutive in further clauses.	Yes. This clarification has been made in the standard.
2	133	Table 5-23	The entry condition in table 5-23 should be *"Lane_Start is asserted and No_Signal is de-asserted"* instead of just *"Lane_Start is asserted".*	The standard has been updated as requested.
3	137	Section 5.9.2.8 c	When three INIT3 control words are received with the same initialisation parameters without intervening RXERR control words, move to the Active state. Assume that side A is entering the Connected state and starting to send INIT3s. After this side B will receive 3 consecutive INIT3s and in turn enter the Connected state. Then side B will receive 3 more consecutive INIT3s and enter the Active state. Note that while being in the Connected state side B has sent only 3 INIT3. Thus, if one or several of these INIT3s transmitted from side B to side A is corrupted, side A will not enter the Active state. In this case side B will be in the Connected state up to the 20 us timer expiration and side A, which is in the Active state, will transmit user data until the retry buffer becomes full or the remote side closes the connection.	Various alternatives to the clock-data recovery lock detection are being considered. The RX_ERROR counter operation will be finalised once this research has been concluded. The idea is that the initialisation state machine will not move to Connected state until the error rate measured is zero or nominal, so the intervening RXERR check will be unnecessary. This has the advantage that initialisation will not complete until the clock-data recovery circuitry is well locked.

#	Page	Section/Figure/Table #	CR	Type
4	139	Section 5.9.2.9 c.4.	The maximum value of the RXERR words counter, which should correspond to the overflow of this counter, is not defined by a certain value. It should be somehow defined, or the recommendations should be given in order to correctly detect a persistent error.	Agreed. Various alternatives to the clock-data recovery lock detection are being considered. The RX_ERROR counter operation will be finalised once this research has been concluded.
5	50	Section 4.4.11	The specification states, that in the LostSync and CheckSync states *"the data and control words being written to the receive elastic buffer are replaced by IDLE control words"*. However, the normative part of the specifications defines that they should be replaced by RXERR control words.	The standard has been updated accordingly.

5.7 Encoding Layer

#	Page	Section/Figure/Table #	CR	Type
1	148	Figure 5-10	There are two comma K-codes which are used in the SpaceFibre. According to the terms and definitions of the SpaceFibre specification *"Comma received"* condition corresponds to the reception of the K28.7 K-code. However, in this case it will be impossible to move from the LostSync state during the initialisation process. All initialisation control words use init commas, which are different from K28.7. We propose to extend the condition to *"comma received OR init comma received"*.	Yes, both commas should be included in the state transition. The standard has been updated accordingly.
2	149	Table 5-32	There is a mistake in Entry statement *"From the CheckSync state, when word realignment occurs or when all symbols within a word are valid."* It should be changed to *"From the CheckSync state,*	The standard has been updated as requested.

#	Page	Section/Figure/Table #	CR	Type
			when word realignment occurs or *<u>when a total of more than four symbols are received that are invalid or contain a disparity error.</u>*"	
3	-	The whole specification	All SpaceFibre control words except INIT1, INIT2 and INIT3 contain K28.7 in the least significant symbol. There are some HS SerDes's which establish symbol alignment on K28.5 only. While after the connection establishment only K28.7 comma symbols are sent over the link, such SerDes's will not be able to maintain symbol alignment.	SpaceFibre will not operate with SerDes devices that perform symbol synchronisation on K28.5 only.

6 Outline Specification

This section provides an update to the outline specification of the proposed SpaceWire-RT protocols.

6.1 Applicable Documents

AD1 Parkes SM, Ferrer Florit A, Gonzalez A, and McClements C, "SpaceFibre", Draft F 0.1, Space Technology Centre, University of Dundee, 6th February 2013.

6.2 SpaceWire-RT Protocol Stack

a) The SpaceWire-RT protocol stack shall be as illustrated in Figure 6-1.

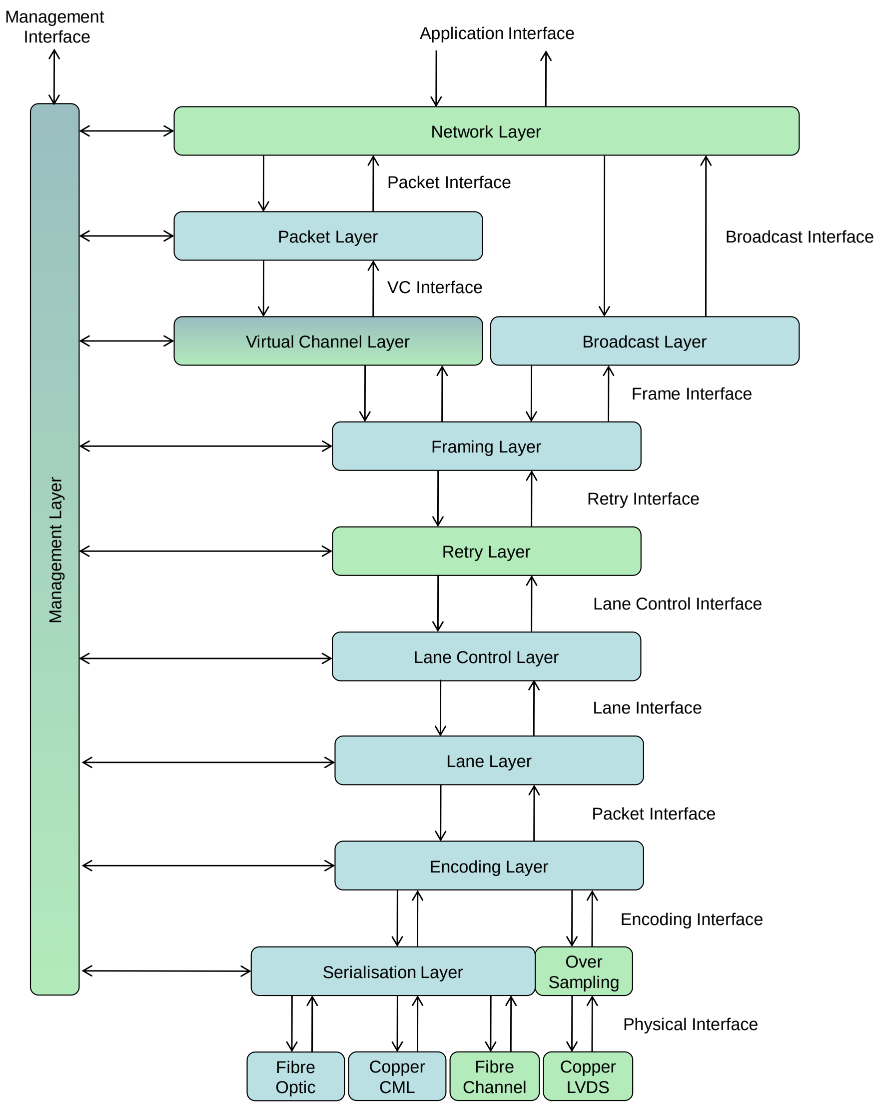

Figure 6-1 SpaceWire-RT Protocol Stack Showing Serialisation and Physical Layer Options

6.3 SpaceWire-RT Application Interfaces

a) There shall be four application interfaces to the SpaceWire-RT protocol stack:

i. SpaceWire Packet Interface

ii. SpaceWire Time-code Interface
iii. Broadcast Message Interface
iv. Management Interface

b) The SpaceWire Packet Interface shall be responsible for sending and receiving SpaceWire packets over a specific virtual channel.
c) The SpaceWire Time-code Interface shall be responsible for sending and receiving SpaceWire time-codes over the SpaceWire-RT network.
d) The Broadcast Message Interface shall be responsible for sending and receiving broadcast messages over the SpaceWire-RT network.
e) The Management Interface shall be responsible for passing configuration, control and status information between the local system and the SpaceFibre device.

6.4 SpaceWire-RT Packets

a) The SpaceFibre Packet Layer shall follow the SpaceFibre packet format defined in the SpaceFibre standard [AD1].

6.5 SpaceWire-RT QoS

a) The Virtual Channel, Framing and Retry Layers shall jointly be responsible for providing quality of service for SpaceWire-RT.
b) The Virtual Channel Layer shall follow the Virtual Channel Layer specification provided in the SpaceFibre standard [AD1].
c) The Framing Layer shall follow the Virtual Channel Layer specification provided in the SpaceFibre standard [AD1].
d) The Retry Layer shall follow the Virtual Channel Layer specification provided in the SpaceFibre standard [AD1].

6.6 SpaceWire-RT Multi-Laning

a) The Multi-Laning layer shall follow the Laning Layer specification provided in the SpaceFibre standard [AD1].

6.7 SpaceWire-RT SpaceFibre Lane

a) The Lane Layer shall follow the Laning Layer specification provided in the SpaceFibre standard [AD1].
b) The Encoding Layer shall follow the Encoding Layer specification provided in the SpaceFibre standard [AD1].
c) The Serialisation Layer shall follow the Serialisation Layer specification provided in the SpaceFibre standard [AD1].
d) The Physical Layer shall follow the Physical Layer specification provided in the SpaceFibre standard [AD1], including both Fibre Optic and Copper media options.

6.8 SpaceWire-RT SpaceFibre with LVDS

a) The Serialisation Layer shall permit the use of oversampling to perform bit synchronisation in the receiver.

   Note: this will significantly reduce the bit rate but will enable implementation without a phase locked loop or similar clock recovery technology.

b) The Physical Layer shall provide an additional option to use LVDS instead of CML running over copper cable.
c) The Physical Layer shall provide an additional option to use Fibre Channel physical layer instead of CML running over copper cable.

6.9 SpaceWire-RT Broadcast Messages

a) The Broadcast Message Layer shall follow the Broadcast Message Layer specification provided in the SpaceFibre standard [AD1].

6.10 Sending SpaceWire Time-Codes as Broadcast Messages

a) SpaceWire time-codes shall be transmitter over a SpaceWire-RT network encapsulated in a broadcast message.

Note: the way in which time-codes are encapsulated in broadcast messages has yet to be defined. One example is to place the time-code time and flags fields into the reserved field of time type of broadcast message. Another possibility is to use a distinct type of broadcast message to carry time-codes.

6.11 Oversampling Serialisation Layer

a) SpaceWire-RT shall permit recovery of the received data stream using oversampling as well as phase-locked loop clock recovery techniques.
b) The two ends of the link shall operate at the same bit rate with a maximum permitted difference in bit clocks between the two ends of the link of 1% (TBC).
c) The receiver bit synchronisation circuitry shall track any change in the receive bit interval and sample the received data bit within +/- 25% of the centre of the bit interval.
d) The received data shall be sampled and de-serialised and passed to the encoding layer for decoding.

6.12 SpaceFibre LVDS

a) SpaceWire-RT shall permit the use of an LVDS physical layer with SpaceFibre.

Note: LVDS is not capable of the Gbits/s signalling speed of CML.

6.13 SpaceFibre Fibre Channel Physical

a) SpaceWire-RT shall permit the use of a Fibre Channel type of physical layer.

Note: This is to provide relatively long distance communication (30 m) at data rates of up to 1 Gbits/s.

7 Conclusion

A coherent set of communication protocols have been defined that covers most of the applications for serial data link and network technology on board spacecraft, including payload data-handling and avionic applications. Research on QoS mechanisms suitable for use with SpaceFibre has been carried out in WP2, resulting in the design of a simple, powerful, and comprehensive quality of service mechanism. This QoS mechanism has then been extended to include specific classes of fault detection in support of FDIR. The QoS mechanisms developed in the SpaceWire-RT project have been adopted into the SpaceFibre standard specification and presented to the SpaceWire working group. In WP3 the SpaceFibre protocols including the QoS and FDIR mechanisms were simulated and the results of the simulations used to help remove errors and inconsistencies and to make necessary clarifications to the SpaceFibre standard, which is appended as annex 1 of this document. The results of the simulation will also be used to inform the VHDL IP Core Development (WP4) and ASIC Feasibility and Prototyping (WP5) activities.

The next stages of the SpaceWire-RT project which are currently running are:

- WP4, VHDL IP Core Development, will explore “oversampling” and SpaceFibre-LVDS, implementing and testing them as an IP core written in VHDL.
- WP5, ASIC Feasibility and Prototyping, will investigate ASIC technologies appropriate for implementation of SpaceWire-RT. Initial design and core prototyping activities will be undertaken, to ensure that the principal risk areas with an ASIC development have been addressed.

8 Annex 1 SpaceFibre Standard

Space
Technology
Centre
University of Dundee

SpaceFibre Specification Draft F 0.3

LIKELY TO CHANGE!

Authors: Steve Parkes
Albert Ferrer
Alberto Gonzalez
Chris McClements

This page is blank intentionally.

Change log

Draft	Authors	Details
Draft A 31st Oct 2007	S.M. Parkes C. McClements	Initial version.
Draft B 22nd Sept 2011	S.M. Parkes	Extensive revisions and re drafted in ECSS format.
Draft C 8th Dec 2011	S.M. Parkes Albert Ferrer	Technical corrections following prototyping and other minor revisions.
Draft D 29th Feb 2012	S.M. Parkes Albert Ferrer Alberto Gonzalez	Technical corrections and clarifications following review.
Draft E1 28th Sept 2012	S.M. Parkes Albert Ferrer Alberto Gonzalez	Retry layer completed. Lane initialisation state diagram simplified. Data word identification state diagram simplified. Control symbols changed to improve robustness. Service interfaces improved. Quality of service specification improved. Management parameters added. Electrical connectors and cable added to physical layer. Other corrections and clarifications throughout the document.
Draft F 0.3 4th Feb 2013	S.M. Parkes Albert Ferrer Alberto Gonzalez	Various updates following simulation of the SpaceFibre standard by SUAI and comments from NEC, MELCO and JAXA.

Table of contents

Figures

Tables

1 Scope

SpaceFibre is a very high-speed serial link designed specifically for use onboard spacecraft. It carries SpaceWire packets over virtual channels and provides a broadcast capability similar to SpaceWire time-codes but offering much more capability. SpaceFibre operates at 10 times the data-rate of SpaceWire and can run over fibre optic or copper media. When operating over fibre optic links SpaceFibre can run over distances of 100 m or more.

SpaceFibre is compatible with the packet level of the SpaceWire standard (ECSS-E-ST-50-12C) and is therefore able to run the SpaceWire protocols defined in ECSS-E-ST-50-51C, 52C and 53C. This means that applications developed for SpaceWire can be readily transferred to SpaceFibre.

The aim of SpaceFibre is to provide point-to-point and networked interconnections for very high data-rate instruments, mass-memory units, processors and other equipment, on board a spacecraft.

This standard covers the protocols required to form a point-to-point link between two units. It does not cover the definition of SpaceFibre packets and SpaceFibre networks, which form the upper layers of SpaceFibre providing compatibility with SpaceWire at those levels.

The SpaceFibre standard specifies the interfaces to the user application and to the physical medium. Some other intermediate interfaces are also specified permitting interoperability at these intermediate levels. The functions that a SpaceFibre CODEC has to implement are specified. Preferred connector and cable characteristics for SpaceFibre optical and copper implementations are also specified.

This standard may be tailored for the specific characteristic and constraints of a space project in conformance with ECSS-S-ST-00.

2
Normative references

The following normative documents contain provisions which, through reference in this text, constitute provisions of this specification. For dated references, subsequent amendments to, or revision of any of these publications do not apply. However, parties to agreements based on this specification are encouraged to investigate the possibility of applying the more recent editions of the normative documents indicated below. For undated references, the latest edition of the publication referred to applies.

ECSS-S-ST-00-01	ECSS system - Glossary of terms
ECSS-E-ST-50-12C	Space engineering - SpaceWire - Links, nodes, routers and networks
ESCC 07072-ST-MDSA HDR-01	ESCC draft specification High data rate connector and assembly “MDSA HDR” for space use.
	Serial ATA Revision 3.0, clause 6.6.1.

3 Terms, definitions and abbreviated terms

3.1 Terms defined in other standards

For the purpose of this Specification, the terms and definitions from ECSS-S-ST-00-01 and ECSS-E-50-50 apply.

3.2 Terms specific to the present standard

NOTE Some rationalisation of the terms will be done in the next revision of this standard.

3.2.1 active lanes

the lanes that are ready to send data and control words, i.e. whose lane initialisation state machine is in the active state

3.2.2 bandwidth credit

the relative amount of link bandwidth that a virtual channel has accumulated

3.2.3 bandwidth credit limit

the maximum amount of positive or negative bandwidth credit that a virtual channel is allowed to accumulate

3.2.4 bandwidth used

the actual portion of link bandwidth that a virtual channel has actually used computed over the previous bandwidth measurement interval

3.2.5 character

data character or control code

3.2.6 comma

K28.7 control code

3.2.7 control code

8B/10B K-code

3.2.8 control word

Comma, K28.0, K28.2 or K28.3 control code followed by three data characters.

3.2.9 current running disparity

the accumulated disparity of a bit stream from when it started to the present moment in time

3.2.10 data character

8-bit data value

3.2.11 data word

word of data comprising four SpaceFibre N-Chars or Fill characters.

3.2.12 disparity

the number of ones in a bit stream minus the number of zeros in that bit stream

3.2.13 even disparity

the same number of ones and zeros in a bit stream

3.2.14 expected bandwidth

the proportion of overall link bandwidth that a virtual channel is expected to use

3.2.15 init comma

initialisation comma

3.2.16 Initialisation comma

K28.5 control code

3.2.17 invalid symbol

symbol that contains a disparity error, i.e. it results in a running disparity greater than one or less than minus one, or is a symbol that does not occur in the 8B/10B decoding table, i.e. is not a valid symbol for an 8-bit data character or control code,

3.2.18 K-code

8B/10B control code

3.2.19 lane

SpaceFibre physical connection between two units

3.2.20 last frame bandwidth

the amount of data sent in the last data frame

3.2.21 link

SpaceFibre connection between two units that incorporates one or more lanes

3.2.22 link bandwidth

the number of data and control words that can be sent over a SpaceFibre link in one second

3.2.23 N-Char

SpaceFibre data character, EOP or EEP

3.2.24 negative disparity

more zeros than ones in a bit stream

3.2.25 neutral disparity

the same number of ones as zeros in a bit stream

3.2.26 Fill

character that can occur in a data frame at the start of a packet to fill the data words containing the destination address and after an EOP or EEP to fill the data word containing the EOP or EEP

3.2.27 ordered set

control word

3.2.28 permanent error

error on a link that cannot be recovered

3.2.29 persistent error

error on a lane that can be recovered only by re-initialising the faulty lane and resending the data

3.2.30 positive disparity

more ones than zeros in a bit stream

3.2.31 priority precedence

the static precedence value of a virtual channel derived from the setting of the priority quality of service management parameter for that virtual channel

3.2.32 ready virtual channel

virtual channel with data ready to send and space in the virtual channel buffer at the far end of the link

3.2.33 required lanes

the lanes that are required to be used to form a SpaceFibre link

3.2.34 reserved bandwidth

the portion of link bandwidth that is set aside for use by a specific virtual channel using the bandwidth reservation quality of service

3.2.35 schedule

list of time slots during which a virtual channel is permitted to send data frames

3.2.36 symbol

10-bit code resulting from 8B/10B encoding of a character

3.2.37 symbol rate

rate at which symbols can be handled in the transmitter and receiver

3.2.38 symbol word

a group of four consecutive symbols that when decoded will form a data word or control word

3.2.39 time slot

an identified interval of time used for scheduling the transmission of data frames

3.2.40 transient error

error on a link that can be recovered by resending the data without re-initialising the link

3.2.41 unrecognised symbol

symbol that does not appear in the 8B/10B symbol table

3.2.42 used bandwidth

the amount of data sent by a particular virtual channel in the last data frame, which is zero for all virtual channels except the one that sent the last data frame

3.2.43 used lane

lane that is being used by the SpaceFibre link

3.2.44 valid symbol

symbol that does not contain a disparity error and is found in the 8B/10B decoding table

3.2.45 word

data word or control word

3.2.46 word rate

rate at which words can be handled in the transmitter and receiver

3.3 Abbreviated terms

The following abbreviations are defined and used within this standard:

Abbreviation	Meaning
8B/10B	8-bit/10-bit
AC	alternating current
ACK	acknowledgement
BC	broadcast channel
BER	bit error rate
B-TYPE	type of data in a broadcast frame
CML	current mode logic
CODEC	coder/decoder
CRC	cyclic redundancy code
DMA	direct memory access
EBF	end broadcast frame
EDF	end data frame
EEP	error end of packet
EOP	end of packet
FCT	flow control token
FDIR	fault detection, isolation and recovery
FIFO	first in first out
FR_SEQ#	frame sequence number
INIT3	initialisation acknowledge control word
ID	identifier
IDLE	idle control word
iLLCW	inverse lane layer control word
Inc	increment
INIT1	initialisation control word
iINIT1	inverse initialisation control word
INIT2	initialisation control word 2
iINIT2	inverse initialisation control word 2
INIT3	initialisation control word 3
Len	length
LLCW	lane layer control word
LOS	loss of signal
LS	least-significant
LSB	least-significant bit

LSYNC	lane synchronisation control word
MAC	medium access controller
MS	most-significant
MSB	most-significant bit
NACK	negative acknowledgement
PCB	printed circuit board
PLL	phase locked loop
PRBS	pseudo-random bit sequence
QoS	quality of service
RMAP	remote memory access protocol
RX	receive
SBF	start of broadcast frame
SDF	start of data frame
SIF	start of idle frame
SOIS	spacecraft onboard interface services
TBA	to be advised
TBC	to be confirmed
TX	transmit
VC	virtual channel
VCB	virtual channel buffer
VHDL	VHSIC hardware description language
VHSIC	very high speed integrated circuit
VML	voltage mode logic

3.4 Conventions

In this document hexadecimal numbers are written with the prefix 0x, for example 0x34 and 0xDF15.

Binary numbers are written with the prefix 0b, for example 0b01001100 and 0b01.

Decimal numbers have no prefix.

A value that is reserved shall be set to zero by the transmitter and should be ignored by the receiver.

4 Principles

4.1 SpaceFibre purpose

The aim of SpaceFibre is to provide point-to-point and networked interconnections for very high data-rate instruments, mass-memory units, processors and other equipment, on board a spacecraft. SpaceFibre is designed to operate using both existing and future space qualified SerDes interface devices, enabling practical implementations of SpaceFibre data-links on board spacecraft.

4.2 Guide to clause 5

Clause 5 of this standard provides the normative requirements. The SpaceFibre specification is separated into several functional layers. Each of these is placed into a separate subsection and provided with a service interface specification and a set of functional requirements. Each requirement has been designed to be verifiable.

Section 5.1 is a short overview of the following sub-sections.

Section ◎ provides the service interface specification for the SpaceFibre interface. There are three service interfaces: the SpaceFibre Packet Service which is used to send and receive SpaceFibre packets over SpaceFibre; the Broadcast Message Service which is used to broadcast and receive short messages with low latency; and the Link Management Service which is used to configure and control the SpaceFibre link and to read status and error information.

Section 5.3 describes the formats of control words, SpaceFibre characters, and frames which are used in SpaceFibre to initialise a link, to send data, and to detect and recover from errors.

Section 5.4 covers the virtual channel layer which is responsible for sending and receiving SpaceFibre packets over SpaceFibre and for providing quality of service. Several qualities of service are supported concurrently by SpaceFibre: best effort, priority, bandwidth reservation, and scheduled. A SpaceFibre packet is sent by placing it into a virtual channel buffer and received by reading it out of the corresponding virtual channel buffer at the other end of the SpaceFibre link. Each virtual channel buffer is configured to provide a specific quality of service. The SpaceFibre packet information is segmented to support interleaving of data from several virtual channels taking into account the quality of service of each virtual channel. The virtual channel layer provides flow control across the link to avoid sending data when there is no room for it in buffers at the far end of the link. A medium access controller in the virtual channel layer is responsible for appropriate multiplexing of data segments over the link, taking into account flow control information and quality of service.

Section 5.6 covers the broadcast message layer which is responsible for broadcasting and receiving short messages with low latency. The broadcast message service is intended to support time broadcast, synchronisation, fault signalling, and event signalling applications. The broadcast message functionality is currently being moved to the Network layer.

Section 5.6 covers the framing layer which is responsible for encapsulating data, broadcast messages and flow control information into frames which are sent and received over the SpaceFibre link. The information in data frames is scrambled to mitigate EMI emissions.

Section 5.8 covers the retry layer which is responsible for error detection, isolation and recovery, at the link level. It adds frame sequence numbers and CRC checksums to the frames and flow control tokens (FCTs) from the framing layer. A retry buffer is provided to hold information until its correct reception at the far end of the link has been acknowledged. If a frame or FCT goes missing or arrives containing an error, the contents of the retry buffer are resent to rapidly recover from the fault. SpaceFibre packet and broadcast messages are delivered without error, which simplifies error handling and FDIR at the application level. Negative acknowledgements are used to support rapid recovery from detected errors. The retry layer also provides a mechanism for sending idle frames when there is no application information to be sent. Idle frames optionally contain a pseudo-random bit sequence which can be used for bit error rate (BER) testing of the link.

Section 5.9 covers the multi-lane layer which is responsible for multi-lane operation of a SpaceFibre link allowing information to be sent over several individual physical lanes to enhance throughput. The way in which multiple lanes are controlled and synchronised is specified, along with the mechanism for distributing information over several lanes on the transmit side and concentrating it back into a single information stream at the receive side of the link. A SpaceFibre link is the logical data link, which can comprise one or more physical lanes. The use of multiple lanes is optional.

Section 5.10 covers the lane layer which is responsible for sending information in the form of a stream of data and control words over a single lane. It provides mechanisms for initialising a lane, re-initialising the link in the event of a persistent error, and adjusting for clock differences between the local clock and clock at the far end of the lane. The lane layer also provides an optional parallel loopback facility for test purposes.

Section 5.11 covers the encoding layer which is responsible for encoding data and control words for sending over a lane and for decoding those received from the other end of the lane. SpaceFibre uses 8B/10B encoding. In the receiver the encoding layer provides 8B/10B symbol synchronisation and data and control word synchronisation. Each data or control word is constructed from four 8B/10B symbols. 8B/10B encoding provides a DC balanced signal which can be AC coupled, supporting galvanic isolation.

Section 5.12 covers the serialisation layer which is responsible for transmitting the 8B/10B symbols as a serial bit stream and for recovering 8B/10B symbols from the received serial bit stream. The receiver provides bit synchronisation to recover the bit stream from the signals received by the physical layer. A mechanism for receive signal inversion is provided to permit freedom in routing the high-speed differential SpaceFibre signals on a PCB. A serial loopback facility is also provided for test purposes.

Section 5.13 covers the electrical and optical physical layers. The electrical characteristics of SpaceFibre drivers, receivers, PCB tracks, connectors and electrical cable are specified. The optical characteristics of the fibre optic version are provided. Where appropriate, connector mechanical information is also provided. SpaceFibre uses current mode logic (CML) for its electrical signalling.

Section 5.14 covers the management layer, which is responsible for configuring and controlling the SpaceFibre interface and for reporting error and status information. The values of the configuration parameters following reset are provided.

Section 5.15 covers conformance of implementation to the SpaceFibre specification and describes permitted partial implementations of the SpaceFibre specification.

4.3 SpaceFibre architecture

An overview of the SpaceFibre CODEC architecture is provided in Figure 4-1.

There are nine conceptual layers to the SpaceFibre CODEC:

- Virtual Channel and Flow Control: responsible for quality of service and flow control over the SpaceFibre link.
- Broadcast: responsible for broadcasting short messages across a SpaceFibre network and for receiving and checking those messages.
- Framing: responsible for framing SpaceFibre packet data, broadcast messages and FCTs to be sent over the SpaceFibre link. It is also responsible for scrambling SpaceFibre packet data for EMI mitigation purposes.
- Retry: responsible for recovering from transient and persistent errors on the SpaceFibre link, and for reporting errors and link failure. Detects missing and out of sequence frames.
- Multi-Lane: responsible for operating several SpaceFibre lanes in parallel to provide a higher data throughput and to provide redundancy with graceful degradation.
- Lane: responsible for initialising the lane, detecting lane errors and re-initialising the lane after an error has been detected.
- Encoding/Decoding: responsible for encoding data into symbols for transmission and decoding symbols into data for reception.
- Serialisation: responsible for serialising and de-serialising SpaceFibre symbols so that they may be transferred over the physical medium.
- Physical: responsible for transferring the electrical signals across a fibre optic or copper medium.

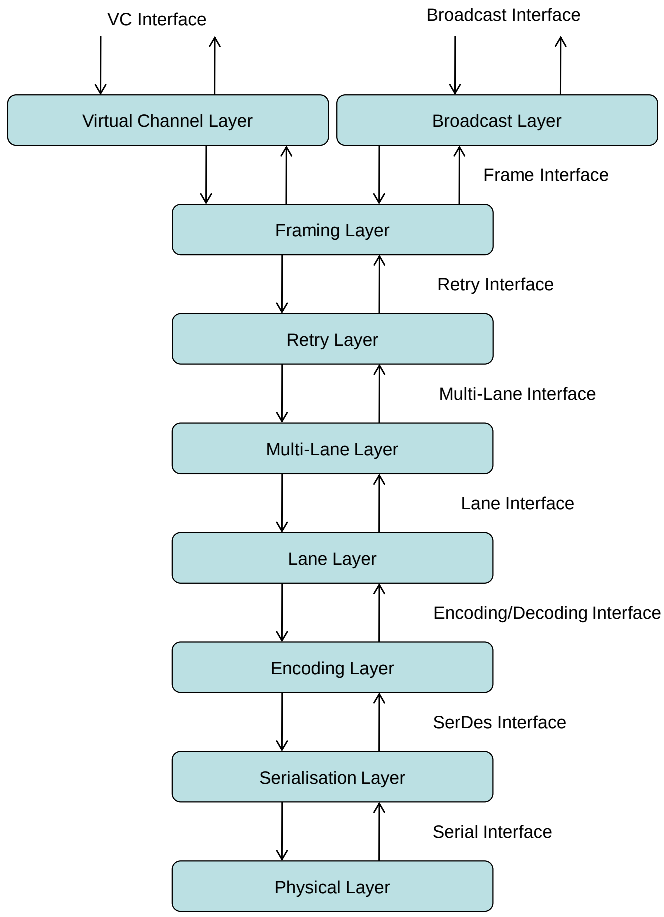

Figure 4-1 Overview of SpaceFibre CODEC

The detailed architecture of the SpaceFibre CODEC is illustrated in Figure 4-2.

The SpaceFibre CODEC performs several functions:

- Virtual channel buffering
- Segmentation
- Flow control
- Quality of service control
- Broadcast sequence number generation
- Broadcast validation
- EMI mitigation
- Framing
- Error rate monitoring
- Frame sequencing
- CRC check
- Frame retry buffering
- Idle frame insertion
- ACK/NACK

- Multi-Lane operation
- Lane distribution/concentration
- Lane synchronisation
- Lane selection
- Lane initialisation and standby management
- Data rate adjustment
- Parallel loop-back
- Word synchronisation
- 8B/10B encoding and decoding
- Symbol synchronisation
- Receiver inversion
- Serialisation and de-serialisation
- Serial loop-back
- Line driver and receiver

Each of these functions is labelled in Figure 4-2 and described in the following sub-sections.

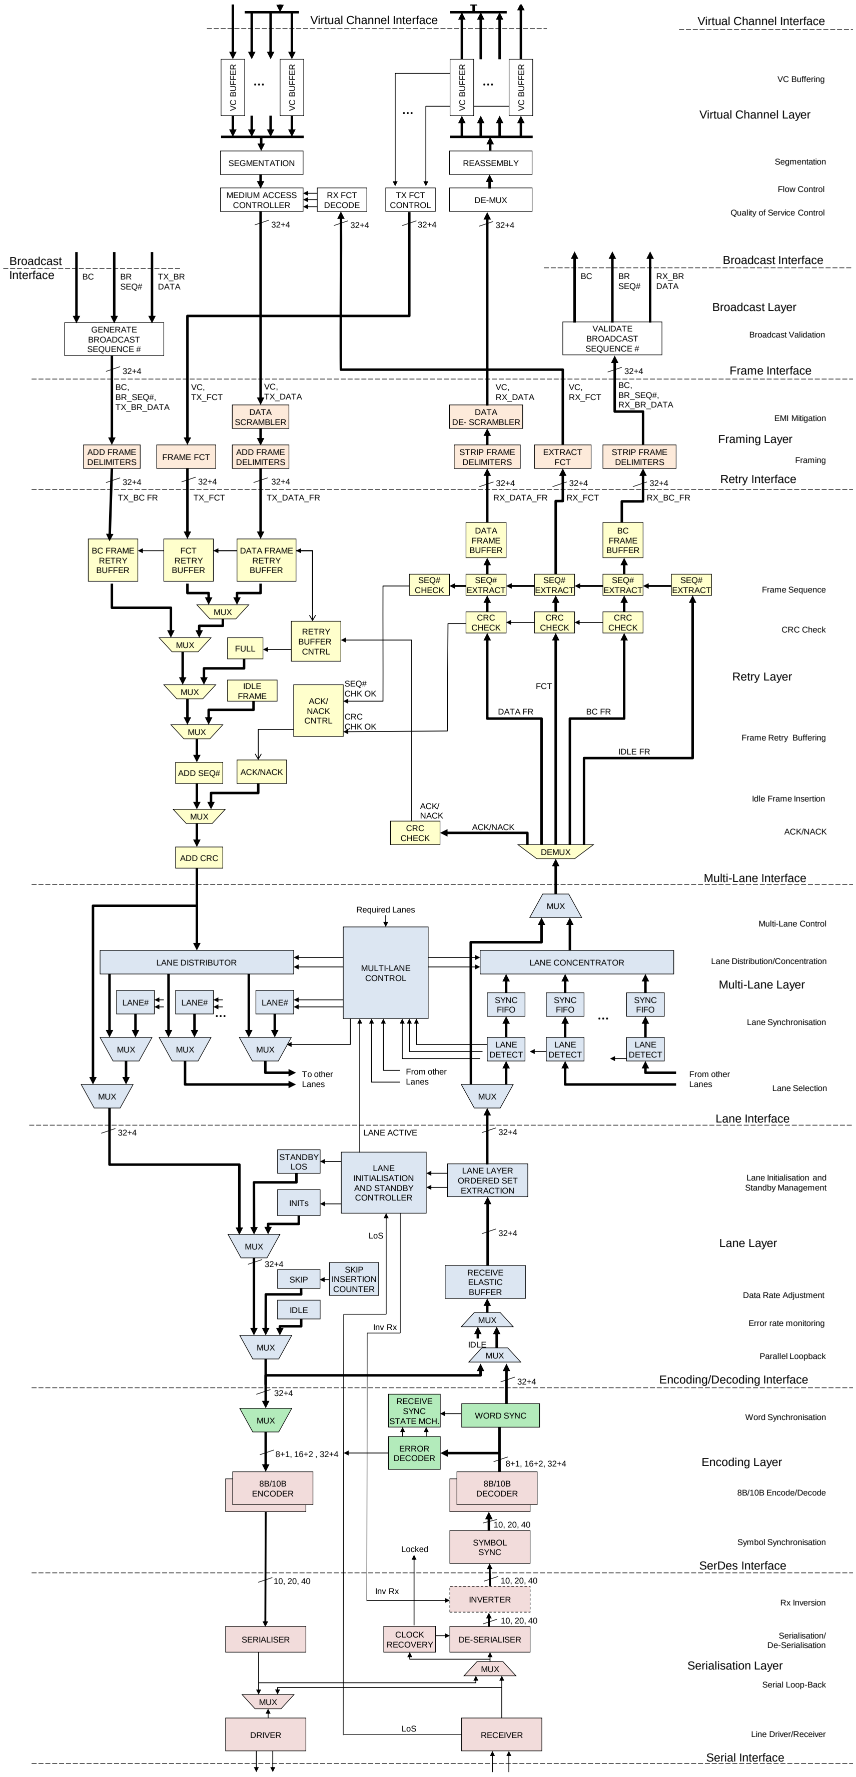

Figure 4-2 SpaceFibre CODEC Conceptual Architecture

4.3.1 SpaceFibre CODEC user interface

There are three different types of interface to the SpaceFibre CODEC: the virtual channel interface used to send and receive SpaceFibre packets, the broadcast channel interface used to broadcast short messages across a SpaceFibre network and to receive those broadcast messages, and the link management interface used to configure and control the SpaceFibre CODEC.

The virtual channel interface of the SpaceFibre CODEC comprises a number of virtual channel buffers for sending SpaceFibre packets (output VC buffers) and the same number for receiving SpaceFibre packets (input VC buffers). There is also an interface for sending broadcast messages and an interface for receiving broadcast messages. The SpaceFibre CODEC is configured and controlled via registers the interface to which is application dependent.

The output VC buffer interface is used to send SpaceFibre packets. Conceptually, each output VC buffer has a FIFO type interface that can accept SpaceFibre data characters and EOP markers. To send a SpaceFibre packet over a SpaceFibre virtual channel, the SpaceFibre packet destination address and cargo are loaded sequentially into the appropriate output VC buffer, followed by an EOP. The specific interface to the VC buffer is application dependent.

Interfaces to the input VC buffers are used to read SpaceFibre packets that have been received over the corresponding SpaceFibre virtual channel. Each input VC buffer has a FIFO type interface, from which SpaceFibre data characters and EOP markers can be read.

The broadcast channel interface to the SpaceFibre CODEC comprises a set of registers for writing the parameters of a broadcast message (broadcast channel, broadcast sequence number, and the message) and a similar set of registers for reading received broadcast messages. The user registers to be notified on the reception of specific classes of broadcast message.

A service interface specification for the SpaceFibre CODEC is provided in section ⊚.

4.3.2 Virtual channel layer

The virtual channel layer is responsible for quality of service and flow control over the SpaceFibre link. It controls the quality of service related to delivery of SpaceFibre packets.

4.3.2.1 Virtual channel buffering

The output virtual channel buffers (VCBs) are used to buffer SpaceFibre packet data before that data is sent over the SpaceFibre link. Data is sent in frames containing up to 256 SpaceFibre N-Chars or Fill characters. The output VCBs permit this amount of data to be buffered before it is offered for transfer over the SpaceFibre link. Sending the data in frames and buffering data prior to framing, permits efficient interleaving of many SpaceFibre packets travelling over different virtual channels over the SpaceFibre link.

The input VCBs provide a similar function for the reception of data arriving over the SpaceFibre interface. An input VCB provides storage for at least one maximum size data frame to ensure that when it arrives there is room for all the data it contains. More space is required to allow data to be transferred continuously if the end user can keep up with the data rate. The application using the SpaceFibre CODEC can then read data from the input VCB at its leisure, without causing loss of data on the SpaceFibre link.

4.3.2.2 Segmentation

Data is sent over the SpaceFibre link in a series of data frames which each contain up to 256 N-Chars or Fill characters. The SpaceFibre packet data in an output virtual channel buffer has to be segmented into chunks of up to 256 N-Chars or Fill characters for placement into the data frames.

When a series of data frames are received for a particular virtual channel the data has to be extracted from them and reassembled into the SpaceFibre packet data which is placed in the appropriate input virtual channel buffer.

4.3.2.3 Flow control

To manage the flow of data from all of the virtual channels across the SpaceFibre link it is necessary to know which output VCBs have data to send at one end of the link, and which input VCBs have space for more data at the other end of the link. Exchange of this information is performed with credit based flow control: by exchanging flow control tokens (FCTs) for data frames. The input VCBs are monitored to determine when they have space for another maximum-sized data frame (up to 256 N-Chars or Fill characters). An FCT is sent to the other end of the link when a particular input VCB has space for another 256 N-Chars or Fill characters. The TX FCT Controller manages the sending of the FCT. When the FCT is received at the other end of the link, it is passed to the RX FCT controller, which informs the MAC which VCBs are able to accept another data frame.

The potential loss of FCTs is handled by the retry layer, which ensures that FCTs cannot be lost unless the link suffers a permanent failure, in which case it is not possible to use that link any more.

4.3.2.4 Quality of service control

A medium access controller determines which VCB is allowed to send data over the SpaceFibre link. This depends on several things:

- Which output VCBs have data to send;
- Which input VCBs at the other end of the SpaceFibre link have space available to receive data;
- The arbitration or quality of service (QoS) policy in force for each Virtual Channel.

For SpaceFibre several quality of service policies are possible including:

- Best effort, where the only Virtual Channels with data to send are those with best effort quality of service, i.e. no other Virtual Channel with a different quality of service has data to send;
- Priority, where the Virtual Channel with the highest priority goes first;

- Bandwidth reserved, where the Virtual Channel with allocated bandwidth and recent low utilisation of the link will go first;
- Scheduled, where time-slots are defined by a local time register or broadcast messages, and only the Virtual Channel allocated to the current time-slot is permitted to send data. If this Virtual Channel has no data to send then another Virtual Channel may use this unused bandwidth opportunistically.

When a virtual channel has data to send in its output VCB and has room for data in the input VCB at the other end of the SpaceFibre link, it competes with other virtual channels in a similar state. The virtual channel permitted to send a frame of data will be the one with the most urgent need to send data according to the QoS policies of all the competing virtual channels. The virtual channel layer then passes a frame containing up to 256 N-Chars or Fill characters from the selected output VCB to the framing layer for sending over the SpaceFibre link.

It is possible to have several virtual channels in a SpaceFibre interface, each operating with a different quality of service. Each virtual channel computes a precedence value based on its quality of service. The medium access controller uses this to determine which virtual channel is permitted to send a frame of data next.

Precedence for the priority quality of service is related to the various priority levels. Extremely urgent priority setting results in a very high precedence value, which will mean that an extremely urgent priority virtual channel will be able to send data just about as soon as it is ready to go. Other priority settings will have lower precedence values and so will compete with other qualities of service when requesting to send data. The best effort quality of service means that the lowest priority is being used.

Precedence for the bandwidth reserved quality of service is determined by the reserved bandwidth for the virtual channel and its recent use of link bandwidth. As a link sends data and uses up its reserved portion of the link bandwidth, so its precedence will drop, making it harder for it to compete with other virtual channels. When it is not able to send data because other virtual channels are sending data, its precedence will increase, making it easier for the link to compete against other virtual channels.

Data flowing in a VC with reserved bandwidth does not cause any other VC to receive less than its reserved bandwidth, provided that the overall bandwidth allocated is less than 100%. A VC is able to use more than its reserved bandwidth, if other VCs are using less than their reserved bandwidth. If one VC is using less than its reserved bandwidth, that spare bandwidth is shared out amongst other VCs.

The scheduled quality of service is based on a schedule table which specifies when a virtual channel is permitted to send data. To support the scheduled quality of service the link bandwidth is split into time-slots, either using a local time register set by a time broadcast message or using a synchronising broadcast message. The schedule table specifies when a virtual channel set to scheduled quality of service is allowed to send data. Such a virtual channel is not allowed to send data at any other time, only in its allocated time-slots. If this virtual channel does not have any data to send, or there is no space in the virtual channel buffer at the other end of the link, another virtual channel is permitted to use the otherwise wasted bandwidth.

All quality of service types can work together in a consistent way, using the multi-layered priority precedence scheme:

- A VC must be scheduled in a particular time-slot to be able to send data.
- VCs with higher priority always have higher precedence.
- Bandwidth reservation applies between VCs set to the same priority level.

A simple scheduled system can be provided by simply assigning VCs to specific time-slots. A simple priority scheme can be implemented by assigning all VCs to every time-slot and giving them individual priority levels. A simple bandwidth reserved scheme can be implemented by assigning all VCs to every-time-slot, giving them equal priority, and reserving an appropriate amount of bandwidth for each VC.

Each virtual channel has its quality of service configured via the link management interface and is able to take on any of the possible SpaceFibre qualities of service.

4.3.3 Broadcast layer

The broadcast layer is responsible for broadcasting short messages across a SpaceFibre network and for receiving and checking those messages.

It receives broadcast messages to be sent from the user application and passes the information necessary to form a broadcast frame to the framing layer.

When a broadcast frame is received from the framing layer, the information it contains is taken out of the frame and passed to the user application.

4.3.3.1 Broadcast messages

A broadcast message is a short message that is sent by a node to all the other nodes on the SpaceFibre network. Broadcast messages propagate in a similar manner to SpaceWire time-codes. Each broadcast message contains a broadcast sequence number which is incremented each time a new broadcast message is sent. When a broadcast message arrives at a SpaceFibre receiver it is checked for errors and its broadcast sequence number is validated by comparing it to the broadcast sequence number of the last broadcast message received. The broadcast message is valid if its broadcast sequence number is one more than that of the previous broadcast message received. Only valid broadcast messages are passed out of the SpaceFibre CODEC. A SpaceFibre router will forward the broadcast message out of all of its SpaceFibre links except the one that the broadcast message was received on.

SpaceWire permits one set of time-codes to be broadcast, although by using the two flags in the time-code it is possible (but not legal according to the SpaceWire standard) to have four independent sequences of time-codes operating concurrently. SpaceFibre broadcast messages permit up to 256 independent sequences of broadcast messages each of which is referred to as a broadcast channel. Each broadcast channel has a broadcast channel identifier and its own broadcast sequence number.

The broadcast channels are split into three types:

- 0-31: Network management broadcast channels.
- 32-253: Node broadcast channels, with each broadcast channel associated with a node that has a logical address of the same value as the broadcast channel number.
- 254 & 255: Reserved broadcast channels.

The network management broadcast channels are split into three sub-types

- 0-3: Time distribution, which are used to provide fault tolerant distribution of system time over the SpaceFibre network.
- 4-7: Synchronisation, which are used to provide synchronisation services over the SpaceFibre network.
- 8-31: Network control, which are used to support configuration, control, and FDIR of a SpaceFibre network.

Broadcast messages also carry 8 bytes of data. A broadcast type field determines the meaning of the 8 bytes of data. For example, when type = TIME, the 8 bytes contain 8 bytes of time information. A broadcast message over one of the time synchronisation channels would typically be of type TIME and the eight data bytes would contain a system time value (un-segmented time).

Typically a particular broadcast channel will be used by a specific node to broadcast information to all other nodes on the SpaceFibre network. This can be used to signal events that occur in that node to other nodes on the network. Different nodes broadcast over different broadcast channels.

A user application of a SpaceFibre CODEC can subscribe to receive broadcast messages from specific broadcast channels and of specific broadcast type. In this way the application will only be notified and receive those broadcast messages that it is interested in.

4.3.3.2 Broadcast sequence number generation

The transmit side of the broadcast layer is responsible for generating the appropriate broadcast sequence number to be included in a broadcast message. Each broadcast channel has its own broadcast sequence number which is incremented each time a broadcast message is sent over that broadcast channel. The broadcast channel number, broadcast sequence number, and message to be sent in the broadcast message are passed to the framing layer for encapsulation into a broadcast frame and transmission over the SpaceFibre interface.

4.3.3.3 Broadcast validation

The receive side of the broadcast layer is responsible for validating the broadcast message by checking its broadcast sequence number. The broadcast layer receives the broadcast channel number, broadcast sequence number and message data from the framing layer for each correctly received broadcast message. It then checks that the broadcast message is valid i.e. its broadcast sequence number is one more than the broadcast sequence number of the last broadcast message received. If this is the case, the broadcast channel number, broadcast sequence number and message data are all passed to the user application. If the broadcast sequence number is not valid that broadcast message is discarded.

Note that the broadcast sequence number is different to the frame sequence number in the retry layer.

4.3.4 Framing layer

The framing layer is responsible for framing SpaceFibre packet data, broadcast messages and FCTs to be sent over the SpaceFibre link. It is also responsible for scrambling SpaceFibre packet data for EMI mitigation purposes.

The framing layer receives SpaceFibre packet and FCT information from the virtual channel layer and broadcast message information from the broadcast message layer and puts this information into frames.

The framing layer passes the resulting data frames, FCTs and broadcast frames to the retry layer. The retry layer controls the way these three streams of information are concentrated into the single stream of data/control words that flows over the SpaceFibre link.

The framing layer receives data frames, FCTs and broadcast frames that have been received and checked for errors by the retry layer. It removes the framing information and descrambles the data in data frames. Data frames and FCTs are passed up to the virtual channel layer and broadcast frames are passed to the broadcast layer.

4.3.4.1 EMI mitigation

Sending a constant bit pattern over a serial link can cause high levels of EM emission due to the energy being concentrated in a few frequency components. To avoid this, data may be scrambled before transmission by convolving the data sequence with a pseudo-random sequence. A pseudo-random sequence approximates white noise and thus has a wide bandwidth. Convolution of the data with a pseudo-random sequence broadens the frequency components, spreading the energy and reducing the peak emission levels.

SpaceFibre uses this technique to mitigate the EM emissions over copper SpaceFibre links. A scrambler is used to scramble the data in each data frame. The original data is then recovered by a de-scrambler at the other end of the link. To help with this process each frame is multiplied by the same pseudo-random sequence i.e. the pseudo-random generator is re-seeded with the same seed at the start of each frame. Broadcast frames and FCTs are not scrambled.

When the link has no other data to send it will send IDLE control words to keep the link active. Sending repeated IDLEs will once again result in excessive EM emission spectral spikes. To avoid this problem, whenever there are no data frames to send an idle frame containing a pseudo-random bit sequence is sent, this is able to reduce the EM emission peaks by 7 to 10 dB. Idle frames are terminated as soon as another data frame becomes ready for transmission, so that the introduction of idle frames does not significantly affect the sending of data frames. Idle frames are generated by the retry layer.

This EMI mitigation technique depends on the spectrum of the data having relatively high peaks which are broadened by convolution with the pseudo-random bit sequence.

The data being scrambled can contain SpaceFibre data characters, EOPs, EEPs and Fill characters. Fills are only used as fillers at the start of a SpaceFibre packet or at the end of a SpaceFibre packet. The data is scrambled but the EOPs, EEPs and Fills are special K-codes and must not be scrambled. The scrambler includes the value of these K-codes in the scrambling of the data, but does not overwrite the K-code with the scrambled data.

4.3.4.2 Framing

Framing is the delimiting of data, broadcast and idle frames by control words that indicate the start and end of the frame. There are three types of start of frame: start of data frame (SDF), start of broadcast frame (SBF), and start of idle frame (SIF). There are two types of end of frame: end of data frame (EDF) and end of broadcast frame (EBF). The idle frame does not have a specific end of frame control word associated with it. The end of idle frame is indicated by the next start of frame (SDF, SBF or SIF) that is received.

4.3.4.2.1 Data frames

A data frame is illustrated in Figure 4-3.

0 7	8 15	16 23	24 31
COMMA	SDF	VC	Reserved
DATA 1 LS	DATA 1	DATA 1	DATA 1 MS
DATA 2 LS	DATA 2	DATA 2	DATA 2 MS
...	...	...	...
DATA N LS	DATA N	DATA N	DATA N MS
EDF	FR_SEQ#	CRC_LS	CRC_MS

Figure 4-3 Data Frame Format

A data frame is a series of 32+4 bit data words delimited by start of data frame (SDF) and end of frame (EDF) control words. Each data word holds four SpaceFibre N-Chars or Fills. The data frame is limited to a maximum frame size of 64 words (256 SpaceFibre N-Chars or Fills). The data field in a data frame can be scrambled when it is transmitted and unscrambled when it is received.

The start of data frame includes the virtual channel number that identifies the virtual channel associated with this data frame.

The end of data frame includes space for information that will be supplied by the retry layer used for checking that there is no missing or duplicate frames (frame sequence number FR_SEQ#) and that there are no errors in the frame (CRC).

Each symbol in the data field of the data frame can contain a SpaceFibre data character or a SpaceFibre EOP or EEP or Fill. A single data frame holds a chunk from a stream of SpaceFibre packets, all of which are associated with the same virtual channel. For example, a data frame could contain any of the following:

- Start of a large SpaceFibre packet,

- Middle of a large SpaceFibre packet,
- End of a large SpaceFibre packet,
- End of a large SpaceFibre packet and start of following packet,
- Several small packets,
- End of a packet, followed by several small packets,
- Etc.

4.3.4.2.2 Broadcast frames

A broadcast frame is illustrated in Figure 4-4.

0 7	8 15	16 3	24 31
COMMA	SBF	BC	B_SEQ#/B_TYPE
DATA 1 LS	DATA 1	DATA 1	DATA 1 MS
DATA 2 LS	DATA 2	DATA 2	DATA 2 MS
EBF	RSVD/LATE	FR_SEQ#	CRC

Figure 4-4 Broadcast Frame Format

A broadcast frame contains the broadcast channel identifier, the broadcast sequence number, a broadcast type and a eight byte message.

The start of broadcast frame includes the broadcast channel number that identifies the broadcast channel (BC) associated with this broadcast frame, along with the broadcast sequence number (B_SEQ#) used to check that the broadcast message is valid.

The broadcast type specifies the type of broadcast message and hence the meaning of the following eight bytes of data.

The end of data frame includes information supplied by the retry layer for checking that there is no missing or duplicate frames (frame sequence number FR_SEQ#) and that there are no errors in the frame (CRC).

The RSVD/LATE field is reserved except for one bit, the LATE bit, which is used to indicate whether the broadcast frame is late due to a retry.

4.3.5 Retry layer

The retry layer is responsible for recovering from transient and persistent errors on the SpaceFibre link. It detects and corrects for missing, out of sequence and duplicate data frames, broadcast frames and FCTs, along with those that contain errors.

The retry layer receives data frames, FCTs and broadcast frames from the framing layer to be sent over the SpaceFibre link.

To support the retry operation it is necessary for the other end of the link to be able to detect missing, out-of-sequence, or erroneous frames and FCTs. This is achieved by adding a sequence number to each frame or FCT and then adding a CRC checksum to the data and broadcast frames. A single common series of sequence numbers is used for data frames, broadcast frames and FCTs, with the incrementing sequence number being applied in the order in which the frames and FCTs are sent over the link.

The CRC checksum covers the complete frame information including the start of frame and the end of frame containing the sequence number.

The complete data frames, broadcast frames and FCTs, are multiplexed together into a single stream of data ready to be sent over the SpaceFibre link. The broadcast frames and FCTs can be inserted at any point within a data frame. While data frames from different virtual channels are interleaved by the MAC in the virtual channel layer, broadcast frames and FCTs are not interleaved but inserted into the stream of data as soon as they are available to be sent. Broadcast frames and FCTs are interleaved, with broadcast frames having priority. FCTs or other control words cannot be inserted in a broadcast frame. The sequence numbers are assigned and added to the frame as the end of the frame passes from the retry layer to the lane layer.

An example showing how data frames, broadcast frames and FCTs are combined to form the data stream is shown in Figure 4-5.

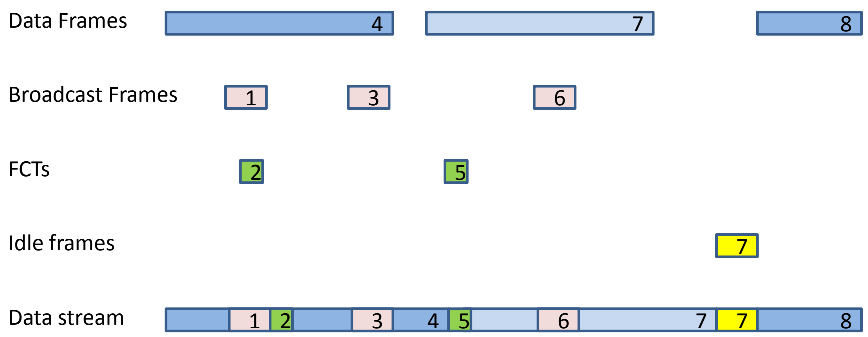

Figure 4-5 Frames and FCTs combined to form the data stream

As the data stream is formed and sent word by word over the SpaceFibre link, it is copied into the frame retry buffers.

On receipt of a data or broadcast frame or FCT over the SpaceFibre link the retry layer checks that it is in the correct sequence and does not contain errors and then passes it up to the framing layer. Each correctly received frame/FCT is acknowledged (ACK) so that the space that it is occupying in the frame retry buffers can be released to make room for other frames waiting to be sent.

When an error is detected a negative acknowledgement (NACK) is sent and current and subsequent received frames/FCTs are discarded until the other end starts a retry process. The NACK contains the frame sequence number of the last correctly received frame/FCT. The erroneous frame/FCT and all subsequent data will be resent together with any new data. The most significant bit of the frame sequence flips every retry process. Hence, the receiver can distinguish between frames/FCTs sent after the retry process and the ones sent before the NACK was received at the other end.

If the SpaceFibre link is re-initialised, the current contents (i.e. not yet ACKed frames/FCTs) of the frame retry buffers are retransmitted. However, broadcast frames and FCTs received after the retry operation can be sent before virtual channel data is retransmitted, as they have higher precedence.

BCs frames that are delayed due to a retry operation are marked as being late (LATE bit set) so that the destination knows that they have been delayed.

When there are no frames/FCTs, ACKs or NACKs to send over the SpaceFibre link the retry layer will generate an idle frame to send. When idle frames are received they are checked to monitor errors and then discarded. Idle frames are given the sequence number of the last sent data frame, broadcast frame or FCT. This enables the receiver to detect a previously missing or discarded frame/FCT and to send a NACK to request its retry. For example consider a data frame that does not arrive for whatever reason. Assume that there is no more data or FCTs to be sent, so an idle frame is sent. This arrives at the far end of the link carrying the sequence number of the data frame that has gone missing. The receiver realises that there is a missing frame and can then send a suitable NACK. Without this facility it could be a long time before there is more data to send prompting the retry. The retry would then rely on a time-out timer, which is much less responsive.

The retry layer sends and receives the data and control words making up the frames/FCTs and ACKs/NACKs to the multi-lane layer

4.3.5.1 Frame sequencing

The retry control function is responsible for checking that frames/FCTs arrive without error (valid words and CRC checksum correct), without duplicate, missing, and out-of-order frames (FR_SEQ# correct). It uses a CRC checksum and frame sequence number to determine if any of these errors have occurred.

There are two sequence number counters: one in the transmitter which produces the sequence numbers and one in the receiver that keeps track of the sequence number of the last correctly received frame/FCT and is used to check the subsequent frame/FCT. On cold reset or remote flush indication, the sequence number in the transmitter and receiver are both set to zero. Remote flush is triggered when the other end of the link is cold reset.

The data frames, broadcast frames and FCTs all share the same frame sequence number. The frame sequence number is not to be confused with the broadcast sequence number used in broadcast frames. Broadcast frames have two sequence numbers: broadcast sequence number and frame sequence number (the latter also being referred to simply as “sequence number”).

When a data frame or broadcast frame or FCT is being sent the sequence number is incremented and its new value placed in the sequence number field.

When a data frame or broadcast frame or FCT is received without error, its sequence number is compared to the value of the receive sequence counter, which holds the sequence number of the last correctly received frame/FCT. If the frame sequence number is one more than the receive sequence counter, that frame is accepted, an ACK sent containing the frame sequence number, and the receive sequence counter incremented.

If the frame sequence number is not one more than the receive sequence counter, the frame/FCT is rejected and a NACK is sent containing the sequence number of the last correctly received frame/FCT, i.e. the current value of the receive sequence counter. The receive sequence counter is not changed.

Frames or FCTs that arrive in error are discarded.

The SIF and FULL control words contains also a sequence number field but the sequence number is not incremented before it is placed in the sequence number field.

When an ACK is received, the space in the retry buffers containing the frames/FCT being acknowledged is freed.

When a NACK is received it contains the sequence number of the last correctly received frame. The space in the retry buffers is freed for all the frames/FCTs up to and including the last correctly received frame. The frames/FCTs that have not been correctly received are then resent, i.e. those with sequence numbers greater than that contained in the NACK.

If one of the frame retry buffers becomes full, due to missing ACKs, FULL words are sent instead of idle frames, which triggers the sending of more ACKs. This will only happen when a system with small retry buffers and very long cables has been implemented. It is intended to indicate a design fault, rather than a run-time error.

4.3.5.2 CRC check

Data frames and broadcast frames contain end user data which has to arrive without error.

A 16-bit CRC checksum is used to detect errors in data frames. This CRC covers the entire data frame including the SDF and EDF control words. The data values of any K-codes are included in the CRC, i.e. the CRC is applied to the 32-bits of the 32+4 bit data stream. Thanks to the error detection of the 8B10B encoding and careful selection of the symbols used in the control words, a single bit-flip error in the bit-stream is detected before this CRC is checked. Therefore a 16-bit CRC error indicates that multiple bit flip errors have occurred within the same data frame which is not expected under nominal BER values. This condition is indicated in a status register.

An 8-bit CRC checksum is used to detect errors in broadcast frames. The CRC covers the entire broadcast frame including the SBF and EBF control words. A shorter CRC can be used since a broadcast frame contains only 16 characters. Broadcast frames are expected to use a small portion of the link bandwidth resulting in lower likelihood of an error in a broadcast frame for a given BER.

An 8-bit CRC checksum is used to detect errors in FCTs, ACKs, NACKs and FULLs. The CRC covers all the fields in these control words.

4.3.5.3 Frame retry buffering

Three frame buffers in the transmit side of the SpaceFibre CODEC are used to store data frames, broadcast frames and FCTs separately until their reception has been acknowledged by the far end of the SpaceFibre link. In the event of a negative acknowledgement (NACK) arriving indicating that one or more frames/FCTs were missing or in error, the retry layer resends all the data frames, broadcast frames and FCTs that have been sent after the one identified by the NACK as being received successfully. The precedence of frames and FCTs are kept: broadcast frames are resent first followed by FCTs and then data frames. Any new broadcast frames and FCTs will be sent before pending data frames are resent.

The frame buffers in the receive side of the SpaceFibre CODEC are used to store a data frame or a broadcast frame while it is being received and checked. Only frames/FCTs without error, without duplication, and in order are passed up to the framing layer for de-framing.

Retry operation is only initiated when a NACK is received. If a link is disconnected and reconnected some frames will be missing and a NACK will be sent. This NACK will initiate the retry of the missing frames. This approach removes the need for a timeout timer.

4.3.5.4 Idle/PRBS

Idle frames are generated by the retry layer when there is no user information to send. On reception, idle frames are checked and then discarded. An idle frame is illustrated in Figure 4-6.

0 7	8 15	16 23	24 31
COMMA	SIF	FR_SEQ#	CRC
SEED LS	SEED	SEED	SEED MS
PRBS 1 LS	PRBS 1	PRBS 1	PRBS 1 MS
PRBS 2 LS	PRBS 2	PRBS 2	PRBS 2 MS
...	...	...	...
PRBS N LS	PRBS N	PRBS N	PRBS N MS

Figure 4-6 Idle Frame Format

An idle frame begins with a start of idle frame (SIF) control word. This begins with a comma, followed by a SIF symbol, a frame sequence number field (FR_SEQ#) that contains the sequence number of the last data frame, broadcast frame, or FCT, and a CRC field which protects the frame sequence number. The frame sequence numbering of an idle frame is a little different to that of a data frame, broadcast frame or FCT: the header of an idle frame contains a frame sequence number, which is the sequence number of the last frame or FCT sent, i.e. the current value of the frame transmit sequence counter is not incremented for an idle frame. An 8-bit CRC is included in the SIF control word to protect the frame sequence number.

The SIF control word is followed by data words that contain a pseudo-random bit sequence (PRBS). The first word is the PRBS word that follows on from the PRBS sequence of the previous idle frame, which forms the PRBS seed for the current idle frame. There is no end of idle frame control word. The idle frame stops as soon as there is any useful data to send, indicated by a start of data frame or start of broadcast frame.

The length of an idle frame is limited to a maximum of 64 PRBS data words. If there is no data frame or broadcast frame to send at that point a new idle frame is started.

4.3.5.5 ACK/NACK

ACKs are sent to indicate that one or more data frames, broadcast frames or FCTs have been received without error and without missing, duplicate, or out of sequence frames or FCTs. An ACK is a control word that contains the frame sequence number of the frame or FCT being acknowledged i.e. of the last frame or FCT correctly received.

NACKs are sent to indicate that a receiving error or a missing data frame, broadcast frame or FCT has been detected. A NACK is a control word that contains the frame sequence number of the last frame or FCT correctly received.

When an ACK is received it causes the frames retry buffers to be updated by removing the frame or FCT that is being acknowledged from its retry buffer. Note that more than one frame or FCT can be acknowledged by a single ACK.

When a NACK is received it causes the contents of the frame retry buffers to be resent, first sending any broadcast frames that need to be resent, then FCTs and then data frames. The NACK also acknowledges any frames/FCTs prior to and including the one it identifies as the last correctly received frame/FCT.

ACKs and NACKs are injected into the transmit data stream with precedence over any data frame, or FCT. A broadcast frame has higher precedence than ACKs and NACKs, to ensure minimum latency and jitter of broadcast frames.

4.3.6 Multi-lane layer

The multi-lane layer permits several lanes to be used in parallel to send data more quickly. For example, if a single lane is capable of sending data at 2Gbits/s, four lanes operating in parallel will permit a data rate of 8 Gbits/s to be achieved. This assumes that the implementation of the higher layers of the SpaceFibre CODEC can support the higher data rate provided by multiple lanes.

The multi-lane layer is responsible for operating several lanes in parallel to provide a higher data throughput and to provide redundancy with graceful degradation.

The multi-lane layer receives data and control words from the retry layer. Depending on the required number of lanes and the number of lanes currently operating, the multi-lane layer spreads the data and control words out over the lanes in use. This is done sequentially on a word by word basis. With the first word going to the first lane in use, the second word to the second and so on.

To ensure that the receiving end of the multilane SpaceFibre link is able to decode the received data and control words in the right order, the multi-lane layer sends lane synchronisation information across each lane whenever the number of available lanes changes. The receiver uses this information to synchronise the decoding of information from the lanes currently in use.

SpaceFibre permits any number of lanes to be operated in parallel up to a maximum of ten lanes (TBC). Since running a link over several lanes reduces the reliability of that link (any one of the N lanes failing will cause the link to fail) SpaceFibre allows additional lanes to be added to support both hot and cold redundancy and to support automatic graceful degradation in the event of a lane failure.

Implementation of lanes is not mandatory. A SpaceFibre CODEC is permitted to have just one lane, in which case it does not need the multi-lane layer.

It is possible for lanes to be crossed over, e.g. lane connector number 1 to be connected to lane connector number 3 at the other end of the link.

4.3.6.1 Multi-lane control

The multi-lane control function is responsible for determining which lanes are to be used to send data/control words and which lanes are receiving data/control words.

On the sending side there are two parameters: the required number of lanes and the number of lanes currently operating (i.e. in the active state, see section 5.10.2). There are then three possibilities:

- Required number of lanes is equal to the number of lanes currently operating, in which case the information to be sent is distributed word by word to the currently operating lanes.
- Required number of lanes is greater than the number of lanes currently operating, in which case the information to be sent is distributed over the available operating lanes. The user application is informed that the requested number of lanes is not available. This provides support for graceful degradation. The link will continue to operate even when one of the requested lanes has failed. The data rate will be reduced but the link continues operating. A network management application could then activate another lane if available, restoring full data-rate operation. This provides cold standby.
- Required number of lanes (N) is less than the number of lanes currently operating, in which case the information to be sent is distributed over the first N lanes. The other operating lanes will only be used in the event of one of the used lanes failing. This provides hot standby.

The multi-lane controller is informed by the lane layer which lanes are active and ready to send information.

4.3.6.2 Lane distribution/concentration

The lane distributor takes the stream of data/control words from the retry layer and spreads them out over the available lanes as determined by the multi-lane controller. The first word goes to the lowest number used lane, the next word to the next lowest number used lane and so on until all the lanes have been provided with a word to send. The next word then goes to the lowest number used lane again. If three lanes are used, words are send in groups of three, one across each of the used lanes.

The lane concentrator does the opposite function to the lane distributor. It takes several data/control words arriving in parallel over the available lanes and passes them sequentially to the retry layer. The lane order in which data/control words arrive is determined by the multi-lane controller.

4.3.6.3 Lane synchronisation

Lane synchronisation is performed following a cold reset of the SpaceFibre interface and whenever the number of required lanes, or number of operating lanes changes. In this event the multi-lane controller causes each operational lane (up to the required number of lanes) to send its lane number over the lane using a lane synchronisation control word. For example, if there are four lanes and the third one is not operating, the other three lanes will be given the lane numbers 1, 2 and 3. Any operating lanes that are not to be used will transfer the lane number 0.

At the other end of the link the synchronisation control word will be detected in each operational lane and the multi-lane controller informed of which lanes are to be used. The multi-lane controller then instructs the lane concentrator which lanes to concentrate data from. Data/control words are then read from the operational lanes sequentially in lane number order.

It is possible that the words arrive in the lanes at different times. This is handled by the lane synchronisation FIFOs. As words arrive they are placed in these FIFOs. They are always read by the lane concentrator in lane number order as specified by the multi-lane controller. A word can only be read from a specific FIFO when it contains one or more data/control words. If a word on a particular lane is late in arriving compared to the other lanes, the lane concentrator is then forced to wait for that word to arrive. Thereafter the lanes will be synchronised.

If a persistent error occurs, the SpaceFibre lane with the error will be re-initialised. While this is occurring, the multi-lane controller will have detected an error and resynchronised the lanes using the lanes that are running properly. Any information that failed to be delivered successfully would then be resent by the retry layer. When the lane with the persistent error has re-initialised, the multi-lane controller will resynchronise the lanes again permitting all lanes to be used once more. If a lane has a permanent error, re-initialisation of the lane will fail. In this case operation with the reduced number of lanes will continue.

It is possible that a lane synchronisation control word is corrupted or lost. If this happens the lane concentrator will not know the correct sequence for reading data from the lanes. There are several possible error cases:

Missing lane synchronisation symbol: detected as an active lane that does not receive a lane synchronisation symbol when the other active lanes receive one. This is likely to be a temporary fault, so all lanes are reinitialised which will automatically cause resynchronisation of the lanes.

Duplicate lane synchronisation symbol: detected as two or more active lanes that receive lane synchronisation symbols with the same lane number. This is likely to be a serious fault in the transmitter and cannot normally occur. All lanes are reinitialised, causing resynchronisation of the lanes to attempt to recover from the fault.

If either type of error persists it must be flagged as a permanent fault to the application.

4.3.6.4 Lane selection

The lane selection function selects whether lanes are to be used or not. On the transmit side a selector determines if the data/control words are to come directly from the retry layer or via the lane circuitry.

On the receive side the data/control words are either passed directly to the retry layer, or via the lane circuitry.

4.3.7 Lane layer

The lane layer is responsible for initialising the lane, detecting link errors and re-initialising the lane after an error has been detected.

Once a lane has been initialised, the lane layer receives data/control words from the multi-lane layer for sending over the lane. It passes data/control words to the encoding/decoding layer for encoding and sending over the lane.

When data/control words are received over the lane, they are passed from the encoding/decoding layer to the lane layer. The lane layer strips out any control words related to the lane layer, using them for lane error rate control, data rate adjustment, lane initialisation and standby management. The data words and remaining control words are passed up to the multi-lane layer.

4.3.7.1 Lane initialisation and standby management

Lane initialisation and standby management is handled by a lane state machine, which is responsible for initialising the lane prior to transfer of data frames, idle frames or broadcast frames. During lane initialisation bit-synchronisation, symbol-synchronisation and word-synchronisation are performed. A handshake protocol is used to ensure that both ends of the lane have achieved synchronisation. Capability negotiation is also performed during initialisation.

The lane initialisation state machine is also responsible for handling Loss of Signal (LoS) indication from the receiver. When many receive errors are detected in a short period in the lane or a Loss of Signal occurs, the lane initialisation state machine sends LOST_SIGNAL control words to the far end of the lane (End A) to indicate that it has suffered LoS. End A of the lane can then record that a Loss of Signal occurred in End B. This is very useful for debugging the cause of a lane failure.

A lane can go into standby when there is no more data to send. In standby the transmitters are disabled to save power while the receivers are enabled ready to auto-start the lane when the other ends tries to initialise the lane again. If one end of the lane (End A) disables its transmitter, it will appear as a Loss of Signal to the other end (End B). To avoid this being detected as an error, End A sends Standby control words to End B to signal that it is about to put the lane into Standby. Both ends are then able to disable their transmitters without causing a Loss of Signal error.

4.3.7.2 Data rate adjustment

The two ends of the lane operate at the same bit rate within the limits of the crystal oscillators at either end. Frequency accuracy is normally ± 100 ppm for a typical crystal oscillator. An elastic-buffer is used in the receiver to compensate for any difference in the clocks at either end of the lane (see Annex A.6 for a description of an elastic buffer). If the transmitter is operating very slightly faster than the receiver then the receive elastic-buffer will start to fill up. If the transmitter is running more slowly than the receiver then the receive elastic-buffer will start to empty. Ideally the receive elastic-buffer should be kept half full. To do this the transmitter sends SKIP control words periodically, at least every 5000 control words or 32-bit data words. If the receive elastic-buffer is more than half full then the SKIP control word is skipped over when reading out of the buffer *i.e.* is ignored and the control word or data following is read out instead. This has then effect of reading out two control words or 32-bit data words rather than one so the buffer becomes emptier. If the receive elastic-buffer is less than half full then the SKIP control word is read out but the buffer output pointer is not updated so that the SKIP is read out twice. This slows down the emptying of the receive elastic-buffer and make the buffer fuller.

The receive elastic-buffer is an important element in the SpaceFibre interface: it allows data or control words to be read out of the 8B/10B encoder every cycle without having to do a hand-shake to see if a character is ready to be read. This significantly improves the speed of operation of the interface.

4.3.7.3 Parallel loop back

A parallel loop back is optionally provided at the bottom of the lane layer. This is used for testing of the upper layers of the SpaceFibre CODEC regardless of the operation of the encoding/decoding, serialisation and physical layers. It is also useful for testing the operation of one end of a SpaceFibre lane before the other end of the lane has been implemented.

4.3.8 Encoding/decoding layer

The encoding/decoding layer is responsible for encoding data/control words into symbols for transmission and decoding symbols into data/control words for reception. It receives data/control words from the lane layer and encodes each of them into four symbols which are then passed to the serialisation layer for serialising and transmission over the physical layer.

Symbols are received from the serialisation layer, but they are not properly synchronised. The encoding/decoding layer synchronises the symbols and decodes them. Groups of four decoded symbols form a data/control word. The encoding/decoding layer is responsible for synchronising the symbols on a data/control word boundary and forming the correctly synchronised data/control words.

4.3.8.1 Word synchronisation

A control word comprises a four symbol code: starting with a comma or other specific K-code (K28.0, K28.2 or K28.3) and followed by three other symbols which are either data-symbols or other K-codes. The comma symbol in a control word is used to indicate the first symbol in a group of four that together form a data/control word. The word synchroniser looks for the comma symbol, re-synchronising the word boundary when a comma is received. This synchronises subsequent data and control words.

4.3.8.2 8B/10B encoding and decoding

The 8B/10B encoding takes an 8-bit input together with a data/control (D/K) flag and provides a 10-bit code for serialisation. The 8B/10B decoder does the inverse operation to recover the 8-bit value and D/K flag. The 8B/10B encoding/decoding is described further in section A.2.

4.3.8.3 Symbol synchronisation

In the receiver once the bit synchronisation has been achieved and the bit stream recovered (see sections 4.3.9 and A.4) it is necessary to determine the boundary of the symbols in the bit stream (every 10-bits) and then to identify the data/control word boundary (every 40-bits). The comma symbol (see section A.2) is used to support both symbol and data/control word boundary detection. The comma symbol contains a unique 7-bit code which can be used to detect the boundary between the comma symbol and the next symbol. Control words normally begin with a comma symbol so the boundary between one data/control word and the next can also be detected using the comma symbol. Note that three control words, flow control token (FCT), end of data frame (EDF) and end of broadcast frame (EBF), do not start with a comma.

4.3.8.4 Synchronisation on single disparity commas

Some radiation tolerant SerDes devices support symbol synchronisation on only positive disparity commas or only negative disparity only commas. This means that there is a possibility that symbol synchronisation might never occur if the symbol stream was such that the commas were always the wrong disparity for synchronisation.

To avoid this situation the control words used during initialisation have been specially designed to ensure symbol synchronisation on positive disparity commas, negative disparity commas, or both.

During initialisation INIT1 control words are sent continuously until the far end of the link has achieved bit, symbol and word synchronisation and sent back INIT2s control words. INIT1s and INIT2s have a special structure to ensure that synchronisation can take place on positive or negative only commas.

The initialisation comma (K28.5) always changes the running disparity e.g. if the disparity is positive immediately before a comma is received, it will be negative immediately afterwards. If the three symbols after the comma in an INIT1 or INIT2 all have even disparity, after they have been received the disparity will remain in the same state as after the comma. So, in the example, if the disparity is positive before an INIT1 is received, it will become negative immediately afterwards. Sending a pair of INIT1s one after the other is therefore enough to ensure proper synchronisation when the symbol synchroniser operates on only positive or only negative disparity commas.

4.3.9 Serialisation layer

The serialisation layer is responsible for serialising and de-serialising SpaceFibre symbols so that they may be transferred over the physical medium. It also provides an optional receiver inversion function which simplifies PCB layout and is recommended for all new SpaceFibre SerDes designs.

The serialisation layer receives symbols from the encoding/decoding layer. These are transferred either singly, in pairs or in groups of four, depending on the specific nature of the serialisation layer (e.g. external SerDes device).

Symbols received over the SpaceFibre physical layer are de-serialised by the serialisation layer and passed up to the encoding/decoding layer either singly, in pairs or in groups of four, depending on the specific nature of the serialisation layer.

4.3.9.1 Receiver inversion

Receiver inversion enables the PCB layout to be optimised for signal integrity. The receiver automatically detects the polarity of the signal being received and corrects the information being received accordingly. This permits the CML + and – signals to be swapped around on both the transmitter and receiver PCBs. The advantage is that the PCB layout can concentrate on signal integrity avoiding having to cross over CML + and – signals using multiple vias which would degrade signal integrity, especially at the high speeds at which SpaceFibre operates.

Some SerDes devices that can be used as part of a SpaceFibre CODEC design do not include receiver inversion and a PCB design using such a SerDes must not swap CML + and – lines. Furthermore if one of these devices is being used transmit signal inversion is not permitted at the far end of the link.

During initialisation the receiver has to detect whether the bit stream is inverted or not. To support this, the INIT1 and INIT2 control word symbols are selected to have valid inverse symbols, i.e. if the symbols is bit-wise inverted it will create a different, but valid symbol. The lane initialisation state machine can check for these inverted INIT1 or INIT2 symbols and change the polarity of the receive bit stream if they are received.

4.3.9.2 Serialisation and de-serialisation

Serialisation is the conversion of the 10-bit parallel symbols to a bit stream.

De-serialisation is the recovery of a 10-bit parallel stream from a serial bit stream. This requires clock recovery from the bit stream using a clock recover circuit which is helped by the 8B/10B code giving a transition-rich set of transmitted codes.

4.3.9.3 Serial loop back

A serial loop-back function is provided to support testing of the complete SpaceFibre CODEC, including lanes, without the need for a physical loop-back. It also permits the testing of the physical layer before the upper layers of the CODEC at the far end of a link have been tested or integrated.

4.3.9.4 Line driver and receiver

The serialised data is driven onto a 100 ohm differential impedance line using current-mode logic (CML). This can then be used to drive a fibre optic transmitter. A fibre optic receiver converts the received optical signals back to CML which is then passed over a 100 ohm differential impedance line to the receiver. Both the CML driver and CML receiver are AC coupled.

CML is used on its own to provide a copper interface to SpaceFibre.

4.3.10 Physical layer

The physical layer is responsible for transferring the electrical signals across a fibre optic or copper medium.

4.4 Typical operation

The typical operation of the SpaceFibre interface will now be described with reference to Figure 4-2.

4.4.1 Sending data over virtual channels

The user application writes SpaceFibre packet data characters and EOPs into one of the output virtual channel buffers (VCBs). The size of the output VCB might be small compared to the size of a packet so the packet information may have to be written in chunks into the output VCB. Each output VCB is large enough to contain at least one (TBC) complete data frames i.e. can contain at least 256 SpaceFibre data characters or EOPs. The user application can write into several of the output VCBs at the same time since multiple SpaceFibre packets can be transferred over the SpaceFibre link concurrently.

When there is enough data in an output VCB to fill a complete data frame, or when the VCB contains at least one EOP, the output VCB informs the medium access controller that it has some information ready to send.

The medium access controller receives information about the space available in each of the input VCBs at the other end of the SpaceFibre link from the RX FCT decoder. When a virtual channel has data to send (output VCB signals that it has data ready to send) and that virtual channel has room in its input VCB at the other end of the SpaceFibre link (RX FCT decoder signals that space is available), the medium access controller puts that virtual channel forward for arbitration.

Each virtual channel is assigned some quality of service (QoS) parameters that determines which of several virtual channels put forward for arbitration will gain access to the SpaceFibre link and be able to send a data frame. Possible QoS include priority, bandwidth reservation and scheduled arbitration. The QoS for a specific virtual channel is configurable.

From the virtual channels put forward for arbitration and the QoS regimes for each of these virtual channels, the medium access controller selects one of the virtual channels and passes data from it to the framing layer. Data frames from different output VCBs are interleaved by the medium access controller, before they are passed to the retry layer.

The RX FCT decoder receives FCTs from the far end of the link. Each FCT gives credit for sending a further 256 N-Chars or Fills. A transmit credit counter contains 256 times the number of FCTs received minus the number of N-Chars or Fills sent, i.e. credit received minus credit spent. Whenever this count is greater than or equal to 256 it indicates that the input VCB at the other end of the link has room to receive another full data frame (64 words).

The input virtual channel buffers at the local end of the link generate FCTs when they have space for another full data frame. These FCTs are produced by the TX_FCT controller which monitors the amount of space in each VCB to determine when it has enough unallocated space for another maximum size data frame. When there is enough space it sends an FCT to the framing layer for transmission to the RX FCT decoder at the other end of the link. The TX FCT controller at one end of the link and the RX FCT decoder at the other end of the link cooperate to control the flow of data frames across the link to avoid overflow of each of the input virtual channel buffers.

When the TX FCT controller needs to send an FCT to the other end of the link it is passed from the virtual channel layer to the framing layer.

4.4.2 Message broadcast

To send a broadcast message the user application simply writes the broadcast type and channel number, and the data to be broadcast to the broadcast message output registers. The user application is permitted to pass broadcast messages for broadcasting over the entire SpaceFibre network at the same time as it is writing SpaceFibre packet data into the output VCBs.

Each broadcast channel has its own broadcast sequence counter. When a broadcast message is to be sent, the sequence counter for that broadcast channel is then incremented, and the broadcast type, channel number, broadcast data and the value of the sequence counter are passed to the framing layer.

4.4.3 Framing

The data sent from the selected virtual channel buffer is segmented into segments containing a maximum of 256 N-Chars or Fills. This data is then scrambled by the data scrambler excluding any EOPs, EEPs or Fills, i.e. only the data characters are scrambled. The scrambled data is then framed by adding a start of data frame control word at the front of the data and an end of data frame control word at the end.

The start of data frame control word contains the following information:

- Comma, which indicates that this is a control word;
- Start of Data Frame, which identifies this control word as a start of data frame;
- Virtual channel number, which identifies the virtual channel that this data frame is being sent over;

The end of data frame control word contains the following information:

- End of Data Frame control word, which identifies this control word as an end of data frame. This control word is K28.0 which is not a comma.
- Space for frame sequence number, which will be filled in by the retry layer;
- Space for a 16-bit CRC, which will be filled in by the retry layer;

FCTs from the virtual channel layer are encapsulated into a single control word containing the following information:

- FCT control symbol, which identifies this control word as an FCT. This control word is K28.3 which is not a comma.
- Virtual channel number, which identifies the virtual channel that this FCT relates to;
- Space for frame sequence number, which will be filled in by the retry layer.
- Space for an 8-bit CRC, which will be filled in by the retry layer;

Broadcast messages are framed by the framing layer but the data is not scrambled.

The start of broadcast frame control word contains the following information:

- Comma, which indicates that this is a control word;
- Start of Broadcast Frame (SBF), which identifies this control word as a start of broadcast frame;
- Broadcast channel number, which identifies the broadcast channel that this broadcast frame is being sent over;
- Broadcast sequence number, which is used for validating the broadcast message before it is propagated by a router or forwarded to an application in a node.
- Broadcast type field, which indicates the type of broadcast message and the meaning of the subsequent eight data bytes.

The end of broadcast frame control word contains the following information:

- End of Broadcast Frame (EBF) symbol, which identifies this control word as an end of broadcast frame. This control word is K28.2 which is not a comma.
- Space for an indication that the broadcast frame has been delayed.
- Space for frame sequence number, which will be filled in by the retry layer;
- Space for an 8-bit CRC, which will be filled in by the retry layer.

The data frames, FCTs and broadcast frames are passed to the retry layer.

4.4.4 Retry

To perform the error detection function for the retry function, it is necessary to add a frame sequence number and CRC checksum to the data frames, broadcast frames and FCT.

The various streams of information that are to be sent over the SpaceFibre link are then multiplexed into a single data stream. Any FCTs that the TX FCT controller requires to be transmitted are implanted amongst the data/control words from data frames. FCTs are thus sent very quickly without having to wait for the end of the current data frame.

Broadcast frames have to be sent as soon as possible since they might be carrying time-critical information. Broadcast frames are implanted amongst the data/control words from data frames, rather than interleaved between data frames. This means that the broadcast frames do not have to wait for the end of a data frame before it can be sent. While an FCT/ACK/NACK/FULL can be implanted inside a data frame it cannot be implanted inside a broadcast frame. Broadcast frames can be implanted at any point in a data frame.

The resulting stream of data and control words containing data frames embedded with FCTs and broadcast frames is passed to the multi-lane layer for sending over the SpaceFibre link. A copy of the data stream is stored into three separate frame retry buffers in case any frame or FCT is corrupted or goes missing.

If there is no information to send over the link, an idle frame is automatically generated to fill the gap between useful pieces of information. This is necessary to keep the link synchronised with low EMI.

In the case of it being necessary to retry sending one or more frames, the frames/FCTs to be resent are read out of the frame retry buffers.

Any ACKs or NACKs that need to be sent to indicate successful reception of a frame/FCT or to indicate that a frame/FCT is missing are implanted into the stream of words being passed to the multi-lane layer.

The precedence of information flowing into the multi-lane layer is therefore as follows:

- Broadcast frames highest precedence
- ACK/NACK
- FCTs
- Data frames
- FULLs
- Idle frames lowest precedence

4.4.5 Multi-lane control - transmit

The transmit multi-lane control function distributes the data and control words from the retry layer across the active lanes that are making up the SpaceFibre link. Each lane has a separate physical connection. Running a SpaceFibre link over several lanes increases the data rate that can be achieved over the SpaceFibre link.

The data and control words from the retry layer are distributed word by word over the used lanes by the lane distributor, starting with the first word going to the lowest number used lane.

To provide a means of collecting the distributed words together again at the other end of the link, the lanes must be synchronised. At the transmit side lane numbers are inserted into each of the active lanes, whenever the multi-lane controller detects that the number of required lanes has changed or one or more of the lanes has re-initialised. These lane numbers identify the order in which data and control words should be read out of the lanes and concentrated in the receiver at the far end of the link.

A final multiplexer in the multi-lane layer selects whether lanes are to be used or not.

The data and control words from the multi-lane layer are passed to the lane layer.

4.4.6 Lane transmission

The lane layer is responsible for establishing a connection over the physical SpaceFibre lane.

During lane initialisation special control words (INIT1, INIT2 and INIT3) are multiplexed into the transmitter to perform the initialisation handshake and configuration under direction of the lane initialisation and standby controller. Once a connection has been established it indicates to the multi-lane control level that the lane is ready (ACTIVE) and is then able to transfer data and control words.

To facilitate putting the lane into standby where the transmitters at both end of the lane are disabled the lane initialisation and standby controller injects standby control words into the transmit data stream. This causes the other end of the lane to disable its transmitter. The local transmitter is then disabled. In standby the receivers are still operational so that each end of the lane is able to respond when requested to re-initialise.

To support FDIR, if there is a very high error rate or the local receiver detects that it has lost the receive signal, i.e. that the signal level is below a minimum acceptable level, the lane initialisation and standby controller will cause Loss of Signal (LoS) control words to be injected into the transmit data stream to inform the other end of the lane of the fault.

To support the receive elastic-buffer, skip control words are injected into the transmit data stream at a rate determined by the Skip Counter and at an interval no longer than 5000 control words or data words (see annex A.6).

Idle control words are sent when the lane is active and there is no other information to send.

The data and control words from the lane layer are passed to the encoder for 8B/10B encoding prior to transmission.

4.4.7 Encoding

The 8B/10B encoder encodes each data character or control code into a 10-bit code, which has special properties to support data transfer and recovery in the receiver (see annex A.2). The data words and control words are four symbols wide so conceptually four cascaded 8B/10B encoders are required to perform the encoding.

4.4.8 Serialisation

The 10-bit code from the 8B/10B encoder is passed to the serialiser which converts it into a serial bit stream. This bit stream is driven on to the communications medium by the driver.

The input to the serialiser can be a 10, 20 or 40-bit wide interface from one, two or four cascaded 8B/10B encoders.

4.4.9 De-serialisation

The receiver takes the incoming bit stream and recovers the bit stream clock using a clock recovery circuit, for example a phase locked loop. The recovered bit clock is used by the de-serialiser to convert the serial bit stream into a slower speed, parallel data stream which can be handled a little more easily. This parallel data stream can be 10, 20, or 40 bits wide depending on the implementation of the SpaceFibre interface.

An inverter in the receiver is able to invert all the bits from the receiver. It does this on command from the lane initialisation and standby controller when an inverted data stream has been detected. The inversion by the inverter will then correct the received data. This capability is useful to simplify PCB layout for the very high-speed SpaceFibre signals. It is optional in SpaceFibre as some space qualified SerDes devices do not support receive bit inversion.

4.4.10 Decoding

The 10-bit symbols are distinguished in the parallel data stream by the symbol synchroniser. This unit determines the position of comma codes in the 10-bit parallel data stream and then outputs the data correctly aligned on the symbol boundaries.

The aligned 10-bit symbols are passed to an 8B/10B decoder which decodes the 10-bit symbols into 8-bit data characters each together with a D/K bit. The 8B/10B decoder is run from the received clock. The 8B/10B decoder detects various forms of error: invalid 10-bit code and incorrect running disparity. If an error occurs this is indicated to the receive synchronisation state machine. An error decoder may be necessary to decode the error indications from the SerDes and 8B/10B decoder into signals that are appropriate for the receive synchronisation state machine.

The stream of data plus D/K-flag characters is passed to the word synchroniser. Each data word starts with a comma symbol in the least significant symbol position. Each control word starts with a comma or other specific K-code depending on the particular control word. The word synchroniser synchronises on the position of the commas in the data stream to be able to extract subsequent data and control words from the symbol stream.

The receive synchronisation state machine monitors the operation of the symbol and word synchronisers and the 8B/10B decoder and determines when the receiver is synchronised or has lost sync.

The received, synchronised data and control words are passed up to the lane layer.

4.4.11 Lane reception

The data or control words are buffered in a receive elastic-buffer, which copes with any differences between the receive clock and the local system clock (see annex A.6). Only words that are correctly synchronised are placed into the receive elastic-buffer. When the encoding/decoding layer indicates that the decoder has lost sync, the data and control words being written to the receive elastic-buffer are replaced by RXERR control words.

The output from the receive elastic-buffer is filtered for SKIP, IDLE and other lane control words. SKIP and IDLE control words are discarded and initialisation control words (INIT1, INIT2 and INIT3) are passed to the lane initialisation and standby controller along with the standby control words (STANDBY and LOS).

The lane initialisation and standby controller provides support for lane initialisation and recovery from errors. During initialisation it ensures that bit and character synchronisations are achieved and that the two ends of the lane are both ready to send and receive data.

The remaining data and control words are passed up to the multi-lane layer.

4.4.12 Multi-Lane control - receive

The multi-lane layer on the receive side concentrates the data and control words from the active lanes into a single data/control word stream that is passed up to the retry layer.

When a lane synchronisation control word is received it is detected by the lane detector in each lane and the lane number that it contains is passed to the multi-lane controller. Lanes that are not to be used for transferring data/control words are indicated with a null lane number. The multi-lane controller instructs the lane concentrator about the order in which data is to be read from the lanes as indicated by the lane number in the lane synchronisation control word received for each lane.

Data and control words are passed into a small synchronisation FIFO. The data in the FIFOs is read out by the lane concentrator in the order determined by the multi-lane controller. The lane with lane number 1 has data read out of its synchronisation FIFO first, followed by the one with lane number 2, and so on until all the indicated lane numbers have been used. Data is not read out of any lane with a null lane number. Any delay on one or more physical lanes causing skew between lanes will be resolved by the ordering of lanes and the synchronisation FIFOs. If data is delayed on one lane the lane concentrator has to wait for that FIFO to have data ready, after going once round the lanes, data will be ready in each lane when the lane contractor wants to read from it. The maximum permitted skew between lanes determines the necessary depth of the synchronisation FIFO, which can be very small.

If lanes are not being used the lane concentrator can be bypassed using a multiplexer.

Data and control words from the lane concentrator are passed to the retry layer.

4.4.13 Error checking

The stream of data and control words are de-multiplexed in the retry layer so that each type of control word, data frame or broadcast frame can be dealt with appropriately.

ACKs and NACKs are stripped out of the data stream, checked and their contents (the frame/FCT number being ACKed or NACKed) is passed to the transmit frame retry buffers. ACKs result in frames/FCTs being deleted from the retry buffers, as they have arrived successfully. NACKs result in frames/FCTs being resent as they were not received correctly.

Data frames have their CRC checked. Any error results in the frame being discarded. If the CRC checksum is correct the frame sequence number is checked and should be one more than the last correctly received sequence number. If this is the case, the sequence checker informs the ACK/NACK controller to send an ACK and the data frame is passed to the framing layer. If the sequence number is not correct, the ACK/NACK controller is asked to send a NACK and the data frame is discarded. While the CRC and sequence number are being checked the data frame is temporarily stored in a data frame buffer. This permits it to be discarded if an error is detected, before the data frame is passed to the framing layer.

The FCT includes a frame sequence number and CRC checksum. The format of the FCT is checked and its frame sequence number and CRC validated. If the frame sequence number and CRC are correct an ACK is requested to be sent and the FCT is passed to the framing layer. If the frame sequence number or CRC are not correct a NACK is requested to be sent and the FCT is discarded.

Broadcast frames are processed in a similar way to data frames with their CRC and frame sequence number being checked.

Any idle frames have their sequence number checked. For an idle frame the sequence number should be the same value as the last received data frame, broadcast frame, or FCT. If this is not the case a NACK is requested to be sent, containing the frame sequence number of the last successfully received frame/FCT.

If the encoding layer detects an error in the decoding, one or more RX_ERROR words will be passed to the lane and retry layer. Any frame being received will then be discarded and the send of a NACK requested if required.

The lane layer will also count the number of RX_ERROR words to monitor the Bit error rate of the lane. Following a leaky bucket algorithm the lane layer will report a persistent error if the number of RX_ERROR words is much higher than expected.

Correctly received data frames, FCTs and broadcast frames are passed up to the framing layer.

4.4.14 Frame stripping

The receive side of the framing layer, strips off the frame to expose the information contained in each frame.

Broadcast channel, broadcast sequence number and broadcast message data are extracted from each broadcast frame and passed up to the broadcast layer.

The FCT with its virtual channel number is passed to the virtual channel layer.

The data within a data frame is unscrambled and passed along with its virtual channel number to the virtual channel layer.

4.4.15 Broadcast message reception

The broadcast message needs to have its sequence number checked against the local broadcast counter for the relevant broadcast channel. It can then be signalled to a local application, or re-broadcast by a router.

4.4.16 Receiving data over virtual channels

The data from the data frame is placed in the appropriate virtual channel buffer as indicated by the virtual channel number in the data frame.

When 256 N-Chars or Fills have been read out of an input virtual channel buffer, there is space for another frame to be received. This is indicated to the other end of the lane by sending an FCT containing the number of the virtual channel that has some more room. This is handled by the TX FCT controller.

Received FCTs are used by the RX FCT decoder to indicate to the medium access controller which virtual channels at the other end of the link have room for another frame.

Data from the input virtual channel buffers can be read out by the user application when required.

4.4.17 Loopback

Two loopback facilities are provided inside the CODEC for test purposes. The first provides loopback of the parallel data and control words and the second loops back the serial data before it is driven onto the physical medium. The serial loopback provides both local and remote loop back, *i.e.* serial data being sent is looped back into the receive de-serialiser and the bit stream data being received is fed back to the line driver.

4.4.18 Reset

There are three types of reset that can be applied to a SpaceFibre interface: cold reset, warm reset and remote flush.

Cold reset returns the device to its power on state.

Warm reset is used after a persistent error has occurred to re-initialise a lane. Warm reset resets the lower layers of the SpaceFibre interface without resetting the upper layers. This permits rapid recovery of errors that may occur on a lane, in conjunction with the retry mechanism. The error recover is thus transparent to the higher layers of the SpaceFibre protocol and hence to the user applications.

Remote reset is used when the other end indicates that the status and the contents of the retry and virtual channel layer should be set to the power on state. This is required to recover from an unexpected power cycle event and maintain data integrity but with some unavoidable data loss.

5
Requirements

5.1 Overview

This section provides the normative requirements for SpaceFibre. It begins, in section ◎ by specifying the services that the SpaceFibre CODEC provides. In section 5.3 the formats of data characters, symbols, words, control words, frames and packets are specified. The subsequent sections specify each of the functional layers of the SpaceFibre CODEC interface:

- Packet layer (To be added)
- Virtual channel layer (section 5.4)
- Broadcast message layer (section 5.6)
- Framing layer (section 5.6)
- Retry layer (section 5.8)
- Multi-lane layer (section 5.9)
- Lane layer (section 5.10)
- Encoding layer (section 5.11)
- Serialisation layer (section 5.12)
- Physical layer (section 5.13)
- Management layer (section 5.14)

5.2 SpaceFibre CODEC service interface specification

5.2.1 SpaceFibre packet service

The service primitives that shall be associated with the SpaceFibre packet service are:

SEND_PACKET_DATA.request;
READ_PACKET_DATA.indication;

5.2.1.1 SEND_PACKET_DATA.request

5.2.1.1.1 Function

The SEND_PACKET_DATA.request primitive shall be used to send SpaceFibre packet data through a virtual channel of a SpaceFibre link.

5.2.1.1.2 Semantics

The SEND_PACKET_DATA.request primitive shall provide the following parameters:

SEND_PACKET_DATA.request (Virtual Channel, SpaceFibre Packet Data)

5.2.1.1.3 When Generated

When the user of the SpaceFibre CODEC has a packet or part of a packet to send over the SpaceFibre link, it shall generate a SEND_PACKET_DATA.request primitive to request to send the SpaceFibre packet over a specific virtual channel of the SpaceFibre link.

5.2.1.1.4 Effect On Receipt

On receipt of the SEND_PACKET_DATA.request primitive the SpaceFibre CODEC shall send the SpaceFibre packet data over the specified virtual channel as soon as permitted by the SpaceFibre medium access controller.

5.2.1.2 READ_PACKET_DATA.indication

5.2.1.2.1 Function

The SpaceFibre CODEC shall pass a READ_PACKET_DATA.indication primitive to the Read SpaceFibre Packet service user to indicate that SpaceFibre packet data has arrived over a particular virtual channel and is waiting to be read.

5.2.1.2.2 Semantics

The READ_PACKET_DATA.indication primitive provides parameters as follows:

READ_PACKET_DATA.indication (Virtual Channel, Data).

5.2.1.2.3 When Generated

The READ_PACKET_DATA.indication primitive shall be passed to the Read Packet service user when SpaceFibre packet data is received.

5.2.1.2.4 Effect On Receipt

The effect on receipt of the READ_PACKET_DATA.indication primitive by the Read Packet service user should result in that service user reading the received SpaceFibre packet data.

5.2.2 Broadcast message service

The service primitives that shall be associated with the broadcast message service are:

BROADCAST_MESSAGE.request;

BROADCAST_MESSAGE.indication;

5.2.2.1 BROADCAST_MESSAGE.request

5.2.2.1.1 Function

The BROADCAST_MESSAGE.request primitive shall be used by the SpaceFibre network layer to requests the SpaceFibre CODEC to send a broadcast message through a SpaceFibre broadcast channel.

NOTE The network layer has not yet been fully defined formally. It will be added to this document or provided as separate SpaceFibre Network Standard document in future.

5.2.2.1.2 Semantics

The BROADCAST_MESSAGE.request primitive shall provide the following parameters:

BROADCAST_MESSAGE.request (Broadcast Channel, Broadcast Sequence Number, Broadcast Type, Late, Message).

5.2.2.1.3 When Generated

When the SpaceFibre network layer has a broadcast message to send it shall generate a BROADCAST_MESSAGE.request primitive to request the SpaceFibre CODEC to send the broadcast message over a specific broadcast channel.

5.2.2.1.4 Effect On Receipt

On receipt of the BROADCAST_MESSAGE.request primitive the SpaceFibre CODEC shall send the broadcast message over the specified broadcast channel immediately, subject to link priority rules.

5.2.2.2 BROADCAST_MESSAGE.indication

5.2.2.2.1 Function

The function of the BROADCAST_MESSAGE.indication primitive shall be to indicate to the SpaceFibre network layer that a broadcast message has arrived over a particular broadcast channel and to pass that message to the SpaceFibre network layer.

5.2.2.2.2 Semantics

The BROADCAST_MESSAGE.indication primitive provides parameters as follows:

BROADCAST_MESSAGE.indication (Application, Broadcast Channel, Broadcast Type, Late, Message).

5.2.2.2.3 When Generated

The BROADCAST_MESSAGE.indication primitive shall be passed to the SpaceFibre network layer, when a valid broadcast message is received.

5.2.2.2.4 Effect On Receipt

The effect on receipt of the BROADCAST_MESSAGE.indication primitive by the SpaceFibre network layer shall be for the network layer to validate and forward valid broadcast messages.

5.2.3 Link management service

The service primitives that shall be associated with the link status service are:

WRITE_MANAGEMENT_PARAMETER.request

READ_MANAGEMENT_PARAMETER.request

READ_MANAGEMENT_PARAMETER.response

LINK_STATUS.indication.

5.2.3.1 WRITE_MANAGEMENT_PARAMETER.request

5.2.3.1.1 Function

The function of the WRITE_MANAGEMENT_PARAMETER.request primitive shall be to write a management parameter to the SpaceFibre CODEC.

5.2.3.1.2 Semantics

a. The WRITE_MANAGEMENT_PARAMETER.request primitive shall provide parameters as follows:

WRITE_MANAGEMENT_PARAMETER.request (Parameter Identifier, New Parameter Value).

5.2.3.1.3 When Generated

The WRITE_MANAGEMENT_PARAMETER.request primitive shall be generated when a management parameter in the SpaceFibre CODEC is to be changed.

5.2.3.1.4 Effect On Receipt

The effect on receipt of the WRITE_MANAGEMENT_PARAMETER.request primitive shall be to update the specified management parameter with the new values provided in the request.

5.2.3.2 READ_MANAGEMENT_PARAMETER.request

5.2.3.2.1 Function

The function of the READ_MANAGEMENT_PARAMETER.request primitive shall be to read the value of a management parameter from the SpaceFibre CODEC.

5.2.3.2.2 Semantics

a. The READ_MANAGEMENT_PARAMETER.request primitive shall provide parameters as follows:

READ_MANAGEMENT_PARAMETER.request (Parameter Identifier).

5.2.3.2.3 When Generated

The READ_MANAGEMENT_PARAMETER.request primitive shall be generated when the value of a management parameter in the SpaceFibre CODEC is to be read.

5.2.3.2.4 Effect On Receipt

The effect on receipt of the READ_MANAGEMENT_PARAMETER.request primitive shall be for the SpaceFibre CODEC to respond with a READ_MANAGEMENT_PARAMETER.response primitive containing the value of the specified management parameter.

5.2.3.3 READ_MANAGEMENT_PARAMETER.response

5.2.3.3.1 Function

The function of the READ_MANAGEMENT_PARAMETER.response primitive shall be to provide the value of a management parameter requested by a READ_MANAGEMENT_PARAMETER.request primitive.

5.2.3.3.2 Semantics

a. The READ_MANAGEMENT_PARAMETER.response primitive shall provide parameters as follows:

READ_MANAGEMENT_PARAMETER.response (Parameter Value).

5.2.3.3.3 When Generated

The READ_MANAGEMENT_PARAMETER.response primitive shall be generated when the SpaceFibre CODEC receives a READ_MANAGEMENT_PARAMETER.request primitive.

5.2.3.3.4 Effect On Receipt

The effect on receipt of the READ_MANAGEMENT_PARAMETER.response primitive by the user of the SpaceFibre CODEC shall be determined by the user application.

5.2.3.4 LINK_STATUS.indication

5.2.3.4.1 Function

The function of the LINK_STATUS.indication primitive shall be to indicate that the status of the SpaceFibre link has changed.

5.2.3.4.2 Semantics

a. The LINK_STATUS.indication primitive shall provide parameters as follows:

   LINK_STATUS.indication(Status).

b. The Status parameter shall contain one of the following status types:

   1. TBD

5.2.3.4.3 When Generated

The LINK_STATUS.indication primitive shall be generated when the status of the SpaceFibre link changes.

5.2.3.4.4 Effect On Receipt

The effect on receipt of the LINK_STATUS.indication primitive by the SpaceFibre CODEC user shall be user defined.

5.3 Formats

In this section the formats of control words and frames are specified.

5.3.1 Control word format

Several types of control word shall be used by SpaceFibre:

- Lane control words
- Lane synchronisation control words
- Retry control words
- Framing control words
- Flow control words
- Receive error indication control words

5.3.1.1 Lane control words

a. The lane control words shall be used to initialise a SpaceFibre lane, to indicate loss of signal, and to indicate that a lane is about to go into standby.

NOTE The lane control words are constructed as shown in Table 5-1.

b. The comma shall be in the least significant symbol position and shall be sent first.

Table 5-1: Lane control words

Name	Control word	Function
SKIP	Comma, LLCW, SKIP, SKIP K28.7, D14.6, D31.3, D31.3	Sent every N control words or data words to support the receiver elastic buffer operation and skip-tick indication. N must be less than or equal to 5000.
IDLE	Comma, LLCW, IDLE, IDLE K28.7, D14.6, D15.6, D15.6	Sent when the link is initialized and the retry layer does not provide valid words to be sent.
INIT1	Init Comma, LLCW, INIT1, INIT1 K28.5, D14.6, D6.2, D6.2	Send as part of the initialisation handshake. D6.2 has even disparity.
inverse INIT1	Init Comma, iLLCW, iINIT1, iINIT1 K28.5, D17.1, D25.5, D25.5	Sent as part of the initialisation handshake if the signals are inverted.
INIT2	Init Comma, LLCW, INIT2, INIT2 K28.5, D14.6, D6.5, D6.5	Send as part of the initialisation handshake. D6.5 has even disparity.
inverse INIT2	Init Comma, iLLCW, iINIT2, iINIT2 K28.5, D17.1, D25.2, D25.2	Sent as part of the initialisation handshake if the signals are inverted.
INIT3	Init Comma, LLCW, INIT3, Capability K28.5, D14.6, D24.1, D0.0-D31.7	Send as part of the initialisation handshake. The capability field describes the capability of the end of the lane sending the INIT3. This can be used to exchange information about the capability of the SpaceFibre interface at the other end of the lane, so that the two ends of the lane can operate in the most efficient way possible.
STAND BY	Comma, LLCW, STBY, STBY K28.7, D14.6, D30.3, D30.3	Indicates that transmitter is moving to the Standby state and will tri-state its driver.
LOS	Comma, LLCW, LoS, LoS K28.7, D14.6, D4.3, D4.3	Indicates that the end of the link sending the

		LOST_SIGNAL control word has lost signal on its receiver.

NOTE The specific values of the K-codes and data symbols used in the control words have been designed to maximise the Hamming distance between one code and other code, helping to reduce the likelihood of an undetected error.

5.3.1.1.2 Skip control word

a. The Skip control word (SKIP) shall be used to support operation of the receive elastic buffer in the SpaceFibre receiver.

b. The Skip control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the Skip control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D14.6 and identifies the control word as being a control word generated and used by the lane layer.

d. The third symbol in the Skip control word shall identify the lane layer control word as being a Skip control word, and has the value D31.3.

e. The fourth and final symbol in the Skip control word shall be a copy of the third symbol.

5.3.1.1.3 Idle control word

a. The Idle control word (IDLE) shall be sent during initialisation and subsequently to keep the lane running when there is no other information to send.

b. The Idle control word (IDLE) shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the Idle control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D14.6 and identifies the control word as being a control word generated and used by the lane layer.

d. The third symbol in the Idle control word shall identify the Lane Layer Control Word as being a Idle control word, and has the value D15.6.

e. The fourth and final symbol in the Idle control word shall be a copy of the third symbol.

5.3.1.1.4 INIT1 control word

a. The INIT1 control word, used during lane initialisation, shall begin with an initialisation comma (K28.5), which is in the least significant symbol position of the control word and is sent first.

NOTE The initialisation comma has the property that its inverse is identical to the non-inverted initialisation comma.

b. The second symbol in the INIT1 control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D14.6 and identifies the control word as being a control word generated and used by the lane layer.

c. The third symbol in the INIT1 control word shall identify the Lane Layer Control Word as being an INIT1 control word, and has the value D6.2.

d. The fourth and final symbol in the INIT1 control word shall be a copy of the third symbol.

5.3.1.1.5 Inverse INIT1 control word

a. The Inverse INIT1 control word (iINIT1) shall begin with an initialisation comma (K28.5), which is in the least significant symbol position of the control word and is sent first.

b. The second symbol in the Inverse INIT1 control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D17.1 and identifies the control word as being a control word generated and used by the lane layer.

c. The third symbol in the Inverse INIT1 control word shall identify the Lane Layer Control Word as being an Inverse INIT1 control word, and has the value D25.5.

d. The fourth and final symbol in the Inverse INIT1 control word shall be a copy of the third symbol.

e. The Inverse INIT1 control word shall not be generated by the SpaceFibre CODEC.

NOTE The Inverse INIT1 is formed when the PCB layout in a SpaceFibre transmitter or receiver crosses over the two signals (CML+ and CML-) making up the differential signal.

5.3.1.1.6 INIT2 control word

a. The INIT2 control word, used during lane initialisation, shall begin with an initialisation comma (K28.5), which is in the least significant symbol position of the control word and is sent first.

b. The second symbol in the INIT2 control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D14.6 and identifies the control word as being a control word generated and used by the lane layer.

c. The third symbol in the INIT2 control word shall identify the Lane Layer Control Word as being an INIT2 control word, and has the value D6.5.

d. The fourth and final symbol in the INIT2 control word shall be a copy of the third symbol.

5.3.1.1.7 Inverse INIT2 control word

a. The Inverse INIT2 control word (iINIT2) shall begin with an initialisation comma (K28.5), which is in the least significant symbol position of the control word and is sent first.

b. The second symbol in the Inverse INIT2 control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D17.1 and identifies the control word as being a control word generated and used by the lane layer.

c. The third symbol in the Inverse INIT2 control word shall identify the Lane Layer Control Word as being an Inverse INIT2 control word, and has the value D25.2.

d. The fourth and final symbol in the Inverse INIT2 control word shall be a copy of the third symbol.

e. The Inverse INIT2 control word shall not be generated by the SpaceFibre CODEC.

NOTE The Inverse INIT2 is formed when the PCB layout in a SpaceFibre transmitter or receiver crosses over the two signals (CML+ and CML-) making up the differential signal.

5.3.1.1.8 INIT3 control word

a. The INIT3 control word, used during lane initialisation, shall begin with an initialisation comma (K28.5), which is in the least significant symbol position of the control word and is sent first.

b. The second symbol in the INIT3 control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D14.6 and identifies the control word as being a control word generated and used by the lane layer.

c. The third symbol in the INIT3 control word shall identify the Lane Layer Control Word as being an INIT3 control word, and has the value D24.1.

d. The fourth and final symbol in the INIT3 control word (Capability) shall contain control flags and information about the capability of the lane, and be a data symbol with any value from D0.0 to D31.7.

NOTE It is not necessary for the Capability field to have valid inverse symbols, since by the time INIT3s are being sent any necessary receiver inversion will have been completed.

e. The Capability symbol shall contain several fields as follows:

1. Bit 0: the Remote_Flush flag,
2. Bit 1: the Lane_Start flag,
3. Bit 2: the Data_Scrambled flag,
4. Bits 3 to 7: reserved, which shall be set to zero when transmitted and ignored when received.

f. The Remote_Flush flag shall take on one of the following values:

0, which means that a flush of the virtual channels and retry layer in the SpaceFibre CODEC receiving the INIT3 control word is NOT required;

1, which means that a flush of the virtual channels and retry layer in the SpaceFibre CODEC receiving the INIT3 control word is required.

g. The Lane_Start flag shall take on one of the following values:

0, which means that the SpaceFibre CODEC sending the INIT 3 control word is set to auto-start;

1, which means that the SpaceFibre CODEC sending the INIT 3 control word is set to lane start.

h. The Data_Scrambled flag shall take on one of the following values:

0, which means that the SpaceFibre CODEC sending the INIT 3 control word will NOT scramble data in its data frames;

1, which means that the SpaceFibre CODEC sending the INIT 3 control word will scramble data in its data frames.

5.3.1.1.9 Standby control word

a. The Standby control word (STANDBY) shall be used to inform the far end of a lane that the SpaceFibre interface is about to go into standby mode with its line driver turned off.

b. The Standby control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the Standby control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D14.6 and identifies the control word as being a control word generated and used by the lane layer.

d. The third symbol in the Standby control word shall identify the Lane Layer Control Word as being a Standby control word, and has the value D30.3.

e. The fourth and final symbol in the Standby control word shall be a copy of the third symbol.

5.3.1.1.10 Loss of signal control word

a. The Loss of Signal control word (LOS) shall be used to inform the far end of a lane that the local receiver is not receiving a signal or that a persistent error has occurred, i.e. too many RX_ERRORs were received in a short period of time.

b. The Loss of Signal control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the Loss of Signal control word shall be the Lane Layer Control Word (LLCW) identifier, which has the value D14.6 and identifies the control word as being a control word generated and used by the lane layer.

d. The third symbol in the Loss of Signal control word shall identify the Lane Layer Control Word as being a Loss of Signal control word, and has the value D4.3.

e. The fourth and final symbol in the Loss of Signal control word (LoS Cause symbol) shall contain information about the cause of the loss of signal, and be a data symbol with any value from D0.0 to D31.7.

f. The LoS_Cause symbol shall contain two fields as follows:

1. Bit 0: the LoS_Cause flag,
2. Bits 1 to 7: reserved, which shall be set to zero when transmitted and ignored when received.

g. The LoS_Cause flag shall take on one of the following values:

0, which means that the receiver in the SpaceFibre CODEC sending the LOS control word is not receiving a strong enough signal ;

1, which means that the receiver in the SpaceFibre CODEC sending the LOS control word is detecting too many receive errors to operate reliably.

h. The Loss of Signal control word shall not be sent when the receiver has received a Standby control word signalling that the other end of the link is about to enter standby mode and turn off its line driver.

5.3.1.2 Lane synchronisation control words

a. Lane synchronisation control words shall be used to synchronise words flowing over multiple lanes.

NOTE Lane synchronisation control words are constructed as illustrated in Table 5-2.

Table 5-2: Lane Synchronisation Control words

Name	Control word	Function
LSYNC	Comma, LSYNC, LANE#, Reserved K28.7, D23.3, D0.0-D10.0, D0.0	Lane Synchronisation. Contains lane number, which indicates the order in which words are to be read from each lane, starting with lane number 1. A lane number of zero (null) indicates that although the lane is running it is not to be used as an active lane. Data will not be sent over a null lane and the lane concentrator will not read data from a null lane.

a. The Lane Sync control word (LSYNC) shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

b. The second symbol in the Lane Sync control word shall identify the control word as being a Lane Sync control word, and has the value D23.3.

c. The third symbol in the Lane Sync control word shall contain lane number, which indicates the order in which words are to be read from each lane, and shall be a data symbol in the range D0.0 to D10.0 (TBC).

d. The lane number D0.0 is the null lane number and indicates that this lane is not being used.

e. The lane numbers D1.0 to D10.0 (TBC) shall indicate the order in which data is to be extracted from each lane, starting with the lowest number lane first.

f. The fourth and final symbol in the Lane Sync control word shall be reserved and set to D0.0.

5.3.1.3 Retry control words

a. Retry control words shall be used to acknowledge data frames, broadcast frames and FCTs that are received correctly and to negatively acknowledge those that are received incorrectly.

NOTE The retry control words are constructed as illustrated in Table 5-3.

Table 5-3: Retry Control word

Name	Control word	Function
ACK	Comma, ACK, FR_SEQ#, CRC K28.7, D2.5, D0.0-D31.7, D0.0-D31.7	Frame Acknowledge. Indicates that a data frame, broadcast frame or FCT has been received without error and in order. Sequence number is the frame sequence number (FR_SEQ#) from the data frame, broadcast frame or FCT that is being acknowledged. CRC is an 8-bit CRC that is used to confirm the integrity of the ACK.
NACK	Comma, NACK, FR_SEQ#, CRC K28.7, D27.5, D0.0-D31.7, D0.0-D31.7	Frame Negative Acknowledge. Indicates that a data frame, broadcast frame or FCT has not been received correctly. Sequence number is the frame sequence number (FR_SEQ#) of the last correctly received data frame, FCT, or broadcast frame. CRC is an 8-bit CRC that is used to confirm the integrity of the NACK.
FULL	Comma, FULL, FR_SEQ#, CRC	Retry buffer full indication.

	K28.7, D15.3, D0.0-D31.7, D0.0-D31.7	Indicates that a retry buffer has become full. To alleviate this situation the other end of the link must send an acknowledgement to a previously received frame/FCT which when received will result in space in the retry buffers being freed. Sequence number (FR_SEQ#) indicates the last data frame, broadcast frame, or FCT sent. CRC is an 8-bit CRC that is used to confirm the integrity of the FULL. FULL will not normally be sent provided that the retry buffers are large enough to handle twice the number of characters that can fit on the line. Hence, FULL will only occur when an implementation has small retry buffers and is operating with a very long cable.
RETRY	Comma, RETRY, Reserved, Reserved K28.7, D7.4, D0.0, D0.0	Retry indication. Indicates to the far end of a link that a NACK has been received and the contents of the retry buffer is about to be transmitted.

5.3.1.3.2 ACK control word

a. The Acknowledgement control word (ACK), shall be used to indicate that a data frame, FCT, or broadcast frame has been received without error and in the correct order.

b. The ACK control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the ACK control word shall identify the control words as being an ACK control word, and has the value D2.5.

d. The third symbol in the ACK control word shall contain a frame sequence number (FR_SEQ#), which is the frame sequence number of the data frame, FCT, or broadcast frame that is being acknowledged, i.e. that has been correctly received at the far end of the link.

e. The fourth and final symbol in the ACK control word shall contain an 8-bit CRC covering all the symbols in the ACK control word, which is used to confirm the integrity of the ACK when received before its contents are acted upon.

NOTE The CRC includes the data part of the comma K-code.

5.3.1.3.3 NACK control word

a. The Negative Acknowledgement control word (NACK), shall be used to indicate that a data frame, FCT, or broadcast frame has not been received correctly.

b. The NACK control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the NACK control word shall identify the control words as being a NACK control word, and has the value D27.5.

d. The third symbol in the NACK control word shall contain a frame sequence number (FR_SEQ#), which is the frame sequence number of the last correctly received data frame, FCT, or broadcast frame.

NOTE All data frames, FCTs, and broadcast frames that have already been sent following that indicated in the NACK will be resent.

e. The fourth and final symbol in the NACK control word shall contain an 8-bit CRC covering all the symbols in the NACK control word, which is used to confirm the integrity of the NACK when received before its contents are acted upon.

NOTE The CRC includes the data part of the comma K-code.

5.3.1.3.4 FULL control word

a. The Full control word (FULL), shall be used to indicate that a retry buffer has become full.

b. The FULL control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the FULL control word shall identify the control words as being a FULL control word, and has the value D15.3.

d. The third symbol in the FULL control word shall contain a frame sequence number (FR_SEQ#) of the last data frame, FCT or broadcast frame sent over the SpaceFibre link.

e. The fourth and final symbol in the FULL control word shall contain an 8-bit CRC covering all the symbols in the FULL control word, which is used to confirm the integrity of the FULL when received before its contents are acted upon.

NOTE The CRC includes the data part of the comma K-code.

5.3.1.3.5 RETRY control word

a. The Retry control word (RETRY), shall be used to indicate that the contents of the retry buffer is about to be sent.

b. The RETRY control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the RETRY control word shall identify the control words as being a RETRY control word, and has the value D7.4.

d. The third and fourth symbols in the RETRY control word are reserved and shall be set to D0.0.

5.3.1.4 Framing control words

a. Framing control words shall be used to encapsulate the data frames, broadcast frames, and idle frames being set across the link.

NOTE Framing control words are illustrated in Table 5-4.

Table 5-4: Data Framing Control words

Name	Control word	Function
SDF	Comma, SDF, VC, Reserved K28.7, D16.2, D0.0-D31.7, D0.0	Start of Data Frame. Contains type of frame and virtual channel, number.
SBF	Comma, SBF, BC, BC_SEQ# K28.7, D29.2, D0.0-D31.7, D0.0-D31.7	Start of Broadcast Frame.
SIF	Comma, SIF, FR_SEQ#, Reserved K28.7, D4.2, D0.0-D31.7, D0.0	Start of Idle Frame. Contains type of frame, and the FR_SEQ# of the last data frame, broadcast frame, or FCT sent. Note there is no end of idle frame control word.
EDF	EDF, FR_SEQ#, CRC_LS, CRC_MS K28.0, D0.0-D31.7, D0.0-D31.7, D0.0-D31.7	End of Data Frame. Contains the frame sequence number and 16-bit CRC for the frame. Note that the EDF starts with K28.0 which is not a comma. This code differentiates all other control words from the EDF control word. Note the sequence number is over the link NOT per VC.
EBF	EBF, RSVD/LATE, FR_SEQ#, CRC K28.2, D0.0, D0.0-D31.7, D0.0-D31.7	End of Broadcast Frame. Contains the frame sequence number and the 8-bit CRC for the frame. Note that the EBF starts with K28.2 which is not a comma. This code differentiates all other control words from the EBF control word. Note the sequence number is over the link NOT per VC.

5.3.1.4.2 Frame sequence number

a. The frame sequence number (FR_SEQ#), used in the Start of Idle Frame (SIF), End of Data Frame (EDF), End of Broadcast Frame (EBF) and in the FCT (see section 5.3.1.5) and FULL (see section 5.3.1.3.4) control words, shall contain two fields: a polarity flag and a 7-bit frame sequence count.

b. The 7-bit frame sequence count field shall contain a modulo-128 integer, which is incremented each time a new data frame, broadcast frame or FCT is sent.

c. The 7-bit frame sequence count field shall be set to zero following a cold reset or remote flush.

d. The polarity flag shall be set to zero following a cold reset or remote flush.

e. The polarity flag shall be inverted every time a new retry is started.

NOTE The polarity flag is used to distinguish frames, ACKs and NACKs sent before a retry starts from those that follow a retry. Each time a new retry starts the polarity bit is flipped to distinguish the new sequence of frames etc. from the old sequence.

5.3.1.4.3 Start of data frame control word

a. The Start of Data Frame control word (SDF), shall be used to indicate the start of a data frame.

b. The SDF control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the SDF control word shall identify the control words as being an SDF control word, and has the value D16.2.

d. The third symbol in the SDF control word shall contain the virtual channel number that this data frame is travelling over.

e. The fourth and final symbol in the SDF control word is reserved and shall be set to D0.0.

5.3.1.4.4 Start of broadcast frame control word

a. The Start of Broadcast Frame control word (SBF), shall be used to indicate the start of a broadcast frame.

b. The SBF control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the SBF control word shall identify the control words as being an SBF control word, and has the value D29.2.

d. The third symbol in the SBF control word shall contain the broadcast channel number that this broadcast frame is travelling over.

e. The fourth and final symbol in the SBF control word shall contain the broadcast sequence number (BC_SEQ#) for the broadcast channel.

NOTE The broadcast sequence number is used to support the broadcasting of the broadcast frame by SpaceFibre routers, in a similar way to the time-code value in SpaceWire supports the broadcast of time-codes.

5.3.1.4.5 Start of idle frame control word

a. The Start of Idle Frame control word (SIF), shall be used to indicate the start of a idle frame.

b. The SIF control word shall begin with a comma (K28.7), which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the SIF control word shall identify the control words as being an SIF control word, and has the value D4.2.

d. The third symbol in the SIF control word shall contain the frame sequence number of the last data frame, FCT or broadcast frame sent over the SpaceFibre link.

e. The fourth and final symbol in the SIF control word shall be reserved and contain the value D0.0.

NOTE There is no end of idle frame control word. Idle frames are ended by a SDF, SBF or SIF.

5.3.1.4.6 End of data frame control word

a. The End of Data Frame control word (EDF), shall be used to indicate the end of a data frame.

b. The EDF control word shall begin with the control code K28.0, which is in the least significant symbol position of the control word and is sent first.

NOTE K28.0 is not a comma.

c. The second symbol in the EDF control word shall contain the frame sequence number of the current data frame.

d. The third symbol of the EDF control word shall contain the least significant byte of a 16-bit CRC which covers the entire data frame including the SDF and EDF.

e. The fourth symbol of the EDF control word shall contain the most significant byte of a 16-bit CRC which covers the entire data frame including the SDF and EDF.

5.3.1.4.7 End of broadcast frame control word

a. The End of Broadcast Frame control word (EBF), shall be used to indicate the end of a broadcast frame.

b. The EBF control word shall begin with the control code K28.2, which is in the least significant symbol position of the control word and is sent first.

NOTE K28.2 is not a comma.

c. The second symbol in the EBF control word shall contain the LATE flag in bit 0, which indicates that the broadcast frame is late, i.e. has been subject to a retry.

d. The remaining bits (bits 1 to 7) of the second symbol of the EBF control word shall be reserved, set to zero when transmitted and ignored when received.

e. The third symbol in the EBF control word shall contain the frame sequence number of the current broadcast frame.

f. The fourth symbol of the EBF control word shall contain an 8-bit CRC covering the entire broadcast frame including the SBF and EBF.

5.3.1.5 Flow control word

a. The Flow control word supports flow control across virtual channels.

NOTE The flow control word is illustrated in Table 5-5.

Table 5-5: Flow control word

Name	Control word	Function
FCT	FCT, Channel#., FR_SEQ#, CRC K28.3, D0.0-D31.7, D0.0-D31.7, D0.0-D31.7	Flow Control Token Indicates that the receive buffer for a specific virtual channel has room for another complete data frame. FCT is a K-code (K28.3) indicating that this control word is an FCT. Channel number identifies the virtual channel which this FCT is for. The FR_SEQ# is a frame sequence number added to the FCT by the retry layer to check for missing, duplicate or out of sequence data frame, broadcast frames and FCTs. CRC is an 8-bit CRC used to ensure that the FCT does not contain any errors.

a. The Flow Control Token control word (FCT), shall be used to indicate that the receive buffer for the specified virtual channel has room for another complete data frame.

b. The FCT control word shall begin with the K28.3 K-code which is in the least significant symbol position of the control word and is sent first.

c. The second symbol in the FCT control word shall contain the virtual channel number that this FCT is for.

d. The third symbol in the FCT control word shall contain the frame sequence number of the current FCT.

NOTE The FCT shares the same frame sequence numbers as the data frames and broadcast frames.

e. The fourth and final symbol in the FCT control word shall contain an 8–bit CRC covering the entire FCT, which is used to check the integrity of the FCT when received, before it is acted upon.

5.3.1.6 Receive error indication control word

a. The receive error indication control word is used by the encoding layer to indicate to a higher layer that a disparity error or invalid code error or other form of error was detected in the received data stream.

NOTE The receive error indication control words are illustrated in Table 5-5.

Table 5-6: Receive error indication control word

Name	Control word	Function
RXERR	Error, Reserved, Reserved, Reserved K0.0, D0.0, D0.0, D0.0	Receive error indication. Indicates that an error has been detected in the received data stream by the decoder. Any word containing one or more symbols in error will be replaced by the RXERR control word. The received data stream is replace by RXERR control words whenever the receive synchronisation state machine is not in the ready state.

b. The receive error indication control word (RXERR), shall be used to indicate that the received data or control word contained an error or is likely to have contained an error.

c. The receive error indication control word shall comprise one error symbol (K0.0) followed by three symbols each set to D0.0.

d. Since the receive error indication control word contains invalid symbols it shall not be transmitted, and is only used in the receiver to indicate receive errors to higher layers.

e. When loss of signal is detected at least one (RXERR) word shall be passed up to the multi-lane or retry layer

f. The RXERR control word may be passed up to the multi-lane or retry layer continuously when the receiver has no signal on its inputs (LOS).

g. Data and control words shall not be passed up to the multi-lane or retry layer when the receiver has no signal on its inputs (LOS).

5.3.2 SpaceFibre Characters

5.3.2.1 SpaceFibre N-Chars

a. SpaceFibre data characters shall be directly represented by data symbols D0.0 to D31.7.

NOTE For example, SpaceFibre data character 0x39 is represented by data symbol D25.1 (see annex A.2).

NOTE SpaceFibre N-chars, data characters, EOPs and EEP are equivalent to SpaceWire N-Chars, data characters, EOPs and EEPs carrying the same type of information, although with a different encoding.

NOTE The SpaceFibre N-Chars are constructed as illustrated in Table 5-7.

Table 5-7: SpaceFibre N-Char Symbols

Name	Symbol	Function
Data	D0.0 to D31.7	Each SpaceFibre character contains one byte of data. Data byte 0x00 is represented by symbol D0.0, data byte 0x01 by D01.0, and so on up to data byte 0xFF which is represented by D31.7.
EOP	K29.7	Represents a SpaceFibre EOP. This can occur at any point in the data field of a data frame, indicating the end of a SpaceFibre packet. The data byte following the EOP is the first byte of the next packet.
EEP	K30.7	Represents a SpaceFibre EEP. This can occur at any point in the data field of a data frame, indicating that the SpaceFibre packet has terminated with an error. The data byte following the EEP is the first byte of the next packet.

b. The SpaceFibre EOP shall be represented by K29.7.

NOTE K29.7 is not a comma.

c. The SpaceFibre EEP shall be represented by K30.7.

NOTE K30.7 is not a comma.

d. One EOP or EEP symbol shall be allowed to appear anywhere within the data field of a data frame.

e. The EOP and EEP symbols shall indicate the end of a SpaceFibre packet.

f. The data symbol following an EOP or EEP symbol shall be the first data symbol of the subsequent SpaceFibre packet.

5.3.2.2 Fill control character

a. The Fill control character shall be represented by K27.7

NOTE K27.7 is not a comma.

NOTE The Fill control character is constructed as illustrated in Table 5-8.

NOTE This character is called a Fill control character because it is used to fill out space at the end of a SpaceFibre packet.

Table 5-8: Fill control word Symbol

Name	Symbol	Function
Fill	K27.7	Used when there is not a multiple of 4 N-Chars in an output VCB to be sent. For example, if there are three N-Chars in an output VCB and the last one is an EOP, it is important to send this tail end of the SpaceFibre packet, without waiting for data from another packet to be added to the buffer. Since SpaceFibre sends data in words of four N-Chars each it is necessary to have a character that will fill out the space at the end of the words. A sequence of Fills will contain one, two or three Fills but never any more.

b. The Fill character shall be allowed to appear anywhere within the data field of a data frame.

NOTE This can be used to support 32-bit alignment of the VCBs. An example showing several small packets in part of a frame, some 32-bit aligned and others packed into 32-bits, is provided in Figure 5-1, where D represents a data character, E an EOP, and F a Fill.

D	D	D	D
D	E	F	F
F	F	D	D
D	D	D	D
E	F	F	F
D	D	E	D
D	D	D	D
E	F	F	F

Figure 5-1 Fills in a virtual channel buffer

c. When the data available from the output virtual channel buffer to be put in the data frame is not a multiple of four N-Chars, Fill characters shall be added to the end of the data field to pad the last word out to a complete data word.

NOTE This will happen when there is one or more EOPs or EEPs in the output VCB. If following the last EOP or EEP added to the output VCB, no more data has been added, there might not be a multiple of 4 N-chars in the buffer. When the last data incorporating the EOP or EEP is read out of the output VCB the space following the EOP or EEP must be filled with Fill characters to form a complete 4 N-Char data word.

d. Fills shall be transported across the SpaceFibre link and placed in the input virtual channel buffer at the receiving end of the link.

5.3.3 Frame Format

a. A frame shall contain user data, from a virtual channel or broadcast channel, or idle data.

b. Three types of frame shall be supported: data frame, broadcast frame, idle frame.

5.3.3.1 Data frame

a. A data frame shall start with a start of data frame (SDF) control word.

NOTE The data frame is illustrated in Figure 5-2.

0 7	8 15	16 23	24 31
COMMA	SDF	VC	RESERVED
DATA 1 LS	DATA 1	DATA 1	DATA 1 MS
DATA 2 LS	DATA 2	DATA 2	DATA 2 MS
...	...	...	...
DATA N LS	DATA N	DATA N	DATA N MS
EDF	FR_SEQ#	CRC_LS	CRC_MS

Figure 5-2 Data Frame Format

b. A data frame shall end with an end of data frame (EDF) control word.

c. A data frame shall contain between one and 64 data words each.

d. Each data word shall contain four SpaceFibre N-Chars.

e. The virtual channel (VC) field in the data frame shall identify the VC sending the data frame and into which it is to be received.

f. The reserved field in the start of data frame (SDF) is reserved and shall be set to D0.0 and shall be ignored by the receiver when the word is identified.

g. The end of data frame shall contain a 16-bit CRC covering the SDF control word, the data in the data frame, and the EDF control word.

h. The frame sequence number (FR_SEQ#) in the end of frame shall contain the frame sequence number for the current frame.

5.3.3.2 Idle frames

a. An idle frame shall start with an idle frame control word.

NOTE An idle frame is illustrated in Figure 5-3.

0 7	8 15	16 23	24 31
COMMA	SIF	FR_SEQ#	CRC
SEED LS	SEED	SEED	SEED MS
PRBS 1 LS	PRBS 1	PRBS 1	PRBS 1 MS
PRBS 2 LS	PRBS 2	PRBS 2	PRBS 2 MS
...	...	...	...
PRBS N LS	PRBS N	PRBS N	PRBS N MS

Figure 5-3 Idle Frame Format

b. An idle frame shall contain between zero and 64 data words.

c. An idle frame shall end with the start of a data frame, broadcast frame, or next idle frame.

NOTE There is no end of idle frame control word.

d. The frame sequence number (FR_SEQ#) shall be the frame sequence number of the last sent data frame, broadcast frame or FCT.

e. The PRBS in an idle frame shall contain a pseudo-random bit sequence (PRBS).

NOTE Sending a PRBS avoids an EMI emission peak when idles frames are being transmitted and enables a PRBS test to be made at the receiver.

f. The Seed field is the first data word following the start of idle frame (SIF) control word and shall be set to a pseudo-random number which is used to seed the PRBS words in the idle frame.

NOTE This permits a SpaceFibre receiver to decode and check the PRBS in the idle frame. This can be used to check the health of the physical layer, after lane initialisation.

g. The PRBS shall be generated using the following algorithm

1. TBD xxx

h. At the end of an idle frame the next pseudo-random word that would have been sent in the next data word of an idle frame shall be kept and used as the seed value of the next data frame.

NOTE In this way the PRBS carries on across a sequence of idle words.

i. The Seed field of the first idle frame following a cold reset shall be set to 0xFFFF.

j. An idle frame shall be terminated as soon as there is a broadcast frame, or data frame to send.

k. An idle frame shall be terminated in any case when 64 PRBS words have been sent.

NOTE This means that the length of the idle frame can contain zero to 64 idle words depending on when there are more broadcast or data frames to send.

l. An idle frame shall contain at least a start of idle frame (SIF) control word.

5.3.3.3 Broadcast frame

a. A broadcast frame shall start with a start of broadcast frame (SBF) control word.

NOTE The broadcast frame is illustrated in Figure 5-4.

0 7	8 15	16 23	24 31
COMMA	SBF	BC	B_SEQ#/B_TYPE
DATA 1 MS	DATA 1	DATA 1	DATA 1 MS
DATA 2 LS	DATA 2	DATA 2	DATA 2 MS
EBF	RSVD/LATE	FR_SEQ#	CRC

Figure 5-4 Broadcast Frame Format

b. A broadcast frame shall end with an end of broadcast frame (EBF) control word.

c. A broadcast frame shall contain two data words each containing four data bytes.

d. The broadcast channel (BC) field in the SBF control word shall identify the broadcast channel transmitting and receiving the broadcast frame.

e. The B_SEQ#/B_TYPE field shall contain two sub-fields: a 3-bit Broadcast Sequence Number (B_SEQ#) field in bit positions 7:5 and a 5-bit Broadcast Type (B_TYPE) field in bit positions 4:0.

f. The Broadcast Sequence Number (B_SEQ#) field in the SBF shall contain an incrementing sequence number specific to the broadcast channel.

NOTE Each BC has its own broadcast sequence number (B_SEQ#) which is used to support the broadcast of the broadcast frame across a SpaceFibre network.

g. The Broadcast Type (B_TYPE) field shall contain a broadcast type, which indicates the type of broadcast message and the meaning of the subsequent eight data bytes.

h. The end of broadcast frame shall contain a RSVD/LATE field which contains two sub-fields: a 7-bit Reserved (RSVD) field and a 1-bit LATE flag.

i. The Reserved (RSVD) field in the EBF shall be set to zero when transmitted and ignored when received.

j. The LATE flag in the EBF shall be set to one when the broadcast frame is resent during a retry, and set to zero otherwise.

NOTE This is used to signal to a receiving node that a broadcast frame has been delayed due to one or more retries. If the broadcast message contains time synchronisation the end user application might decide to ignore the broadcast message because it was late. If the broadcast message contains event signalling information it might still be useful to the application even though it arrives late.

k. The frame sequence number (FR_SEQ#) in the end of broadcast frame shall contain the sequence number of the frame.

NOTE This sequence number is used to support retry.

l. The end of broadcast frame shall contain an 8-bit CRC covering the SBF control word, the data in the data frame, and the EBF control word.

5.3.4 Control word and frame precedence

a. During lane initialisation, prior to a link being established, only the following Lane Layered Control Words shall be sent: SKIP, INIT1, iINIT1, INIT2, iINIT2, INIT3, IDLE.

b. During lane initialisation the Lane Layered Control Words shall have the following precedence:

(1) SKIP, highest precedence

(2) INIT1, iINIT1, INIT2, iINIT2,INIT3.

(3) IDLE, lowest precedence.

c. Once one or more lanes have been established to form a link, the control words and frames shall have the following precedence:

(1) SKIP

(2) LOST_SIGNAL

(3) Standby

(4) LSYNC

(5) RETRY

(6) Broadcast frame

(7) ACK/NACK

(8) FCT

(9) Data frame

(10) FULL

(11) Idle frame

(12) IDLE, lowest precedence.

NOTE It is not possible for an ACK and NACK to be waiting to be sent at the same time: one replaces the other.

d. A SKIP control word shall be inserted within a data frame, broadcast frame or idle frame.

e. A LOST_SIGNAL control word shall be inserted within a broadcast, data or idle frame.

f. A Standby control word shall be inserted within a broadcast, data or idle frame.

g. An LSYNC control word shall be inserted within a broadcast, data or idle frame.

h. A RETRY control word shall be inserted within a broadcast, data or idle frame.

i. A broadcast frame shall be inserted within a data.

j. A broadcast frame shall end an idle frame.

k. An ACK control word shall be inserted within a data or idle frame.

l. A NACK control word shall be inserted within a data or idle frame.

m. An FCT control word shall be inserted within a data frame.

n. An FCT control word shall end an idle fame.

o. An FULL control word shall be inserted within a data or idle frame.

p. A data frame shall end an idle frame.

q. An idle frame shall be sent only when there are no data frames, broadcast frames or FCTs to send.

r. An IDLE control word shall be sent only when there are no other data or control words to send.

s. While the contents of the retry buffer are being resent (see section 5.8), new data frames, FCTs, or broadcast frames cannot be sent and have to wait for the retry buffer to finish resending.

5.3.5 K-code summary

The meaning of K-codes within SpaceFibre shall be as summarised in Table 5-9.

Table 5-9: Meaning of K-codes

K-code	Meaning	Disparity
K0.0	Rx Error	-
K28.0	EDF	even
K28.1	Not used	+2 or -2
K28.2	EBF	+2 or -2
K28.3	FCT	+2 or -2
K28.4	Not used	even
K28.5	Initialisation Comma	+2 or -2
K28.6	Not used	+2 or -2
K28.7	Comma	even
K23.7	Not used	even
K27.7	SpaceFibre Fill	even
K29.7	SpaceFibre EOP	even
K30.7	SpaceFibre EEP	even

5.3.6 Control word symbol summary

The data values assigned to particular control word symbols within SpaceFibre shall be as summarised in Table 5-9.

Table 5-10: Meaning of control word symbols		
D-code	**Meaning**	**Disparity**
D14.6 / D17.1	LLCW/i_LLCW	even
D23.3	LSYNC	odd
D2.5	ACK	odd
D27.5	NACK	odd
D7.4	RETRY	odd
D16.2	SDF	odd
D29.2	SBF	odd
D4.2	SIF	odd
D15.3	FULL	odd
D31.3	SKIP	odd
D15.6	IDLE	odd
D6.2 / D25.5	INIT1/iINIT1	even
D6.5 / D25.2	INIT2/iINIT2	even
D24.1	INIT3	odd
D30.3	STANDBY	odd
D4.3	LoS	odd

5.4 SpaceFibre packet layer

The SpaceFibre packet layer has yet to be written.

SpaceFibre packets have identical form to SpaceWire packets.

5.5 Virtual channel layer

5.5.1 Virtual channel buffering

a. The application interface to the SpaceFibre CODEC shall comprise one or more output virtual channel buffers and the same number of input virtual channel buffers.

b. There shall be a maximum of 256 virtual channels.

c. There shall be an output virtual channel buffer and an input virtual channel buffer for each virtual channel.

d. A particular virtual channel shall be identified by its virtual channel number.

e. An output virtual channel buffer shall be large enough to hold at least 256 SpaceFibre N-Chars.

NOTE This corresponds to one full SpaceFibre data frame.

f. To send a SpaceFibre packet over a particular virtual channel the application shall write the N-Chars forming the SpaceFibre packet into the virtual channel buffer, starting with the leading N-Chars of the SpaceFibre packet.

g. If the output virtual channel buffer becomes full the application shall wait until there is more room in the output virtual channel buffer.

NOTE SpaceFibre packets can be any length. Packets greater than the size of the output virtual channel buffer will not fit into the buffer and will have to be written in as space is made available in the buffer when data is read out and sent over the SpaceFibre link.

h. As soon as one SpaceFibre packet has been written into the output virtual channel buffer, it shall be possible to start writing the next SpaceFibre packet to be sent over that virtual channel, provided that there is room in the output virtual channel buffer for at least one more N-Char.

NOTE This means that it is possible to have many small packets waiting for transmission in the output virtual channel buffer.

i. SpaceFibre packets shall be sent over a virtual channel in the order in which they are written into the output virtual channel buffer for that virtual channel.

j. An output virtual channel at one end of a SpaceFibre link shall be paired with the input virtual channel at the other end of the SpaceFibre link which has the same virtual channel number.

k. When SpaceFibre packets are received they shall be placed in the input virtual channel buffer for the virtual channel that they were sent over.

l. An input virtual channel buffer shall be large enough to hold at least 256 SpaceFibre N-Chars.

NOTE This value is implementation dependent, but normally the buffer size would be sufficient for at least 1024 SpaceFibre N-Chars.

m. SpaceFibre N-Chars shall appear in an input virtual channel buffer in the same order in which they were written into the corresponding output virtual channel buffer at the other end of the link.

n. When there are N-Chars available in an input virtual channel buffer, the host application shall be informed that there is data available to read from that input virtual channel buffer.

o. The host application can read data from the input virtual channel buffer whenever it wants to, provided that there is data available in the input channel buffer.

p. When during initialisation the INIT3 capability field requests a remote flush, all output virtual buffers shall be flushed and for each output virtual buffer, if the last character written by the host application was not an EOP, all new data characters will be discarded up to and including the next EOP.

NOTE This is similar to SpaceWire error handling, which spills the current packet being transmitted when a disconnection error occurs. This ensures that the receiver will not get an incomplete packet with an invalid header.

q. When during initialisation the INIT3 capability field requests a remote flush, all input virtual buffers shall be flushed and for each input virtual buffer, an EEP shall be the next symbol read by the host application if the last symbol read was not an EOP.

NOTE This is equivalent to SpaceWire error handling, which uses the EEP to indicate at the receiver side that the link was disconnected and that the last chuck of packet provided is in error.

5.5.2 Segmentation

a. Each output virtual channel buffer shall keep track of the number of N-Chars written into it and the number read out.

b. When there are at least 256 N-Chars in the output virtual channel buffer, or it contains at least one EOP or EEP, that virtual channel buffer shall indicate that is has data ready to form a data frame to a medium access controller, which controls which output channel is to be permitted to next send a data frame over the SpaceFibre link.

c. When informed that it is now permitted to send a data frame over the SpaceFibre link, an output virtual channel buffer shall pass up to 256 N-Chars, to the medium access controller for sending over the SpaceFibre link.

d. If the output virtual channel buffer contains one or more EOPs or EEPs and less than 256 N-Chars, it shall send all the N-chars it contains to the medium access controller.

e. If the number of N-Chars being sent in a data frame is not a multiple of four N-Chars, one, two or three Fills shall be added to pad out the last data word of the data frame.

f. When a data frame is received, the N-Chars it contains shall be placed in the appropriate input virtual channel buffer.

5.5.3 Flow control

5.5.3.1 TX FCT Control

a. Each input virtual channel buffer shall keep track of the number of N-Chars and Fills written into it from the SpaceFibre receiver, and read out using an input space counter, for that input virtual channel buffer.

NOTE If a frame contains an EOP or EEP it is possible that as well as N-Chars it contains some Fills which are used to pad out the number of N-Chars to a multiple of four. At the link level these Fills are treated in the same way as N-Chars and are transferred across the link and placed in the input virtual channel buffer at the far end of the link.

b. On cold-reset or remote flush, the input space counter shall be set to the amount of space in the buffer, counted as the number of N-Chars and Fills the buffer can contain.

c. On warm-reset, the input space counter shall not be modified.

d. If the value of the input space counter is greater than or equal to 256 the input VC buffer shall request the transmitter to send an FCT.

NOTE After cold reset this can result in several FCTs being sent for a virtual channel if its input buffer can hold more than 256 N-Chars or Fills.

e. For each FCT sent by a specific input virtual channel, 256 shall be subtracted from the corresponding input space counter to reserve that much space for the N-Chars and Fills that the FCT will permit to be sent from the far end of the link.

NOTE An FCT is effectively exchanged for 256 N-Chars and Fills.

f. Each time the end user application reads out an N-Char or Fill from the input VC buffer, the input space counter shall be incremented by one.

g. The input space counter shall indicate the amount of space available in the input virtual channel buffer.

h. When FCTs are being requested to be sent for several different virtual channel buffers, they shall be arbitrated fairly.

i. The input space counter shall have a minimum size of 256 N-Chars or Fills.

NOTE Normally the input VC buffers and hence input space counter would be significantly larger than 256.

5.5.3.2 RX FCT Control

a. Each output virtual channel buffer shall contain a FCT credit counter which shall indicate how much more data it is permitted to send.

b. On cold-reset or remote flush, the FCT credit counter shall be set to zero.

c. On warm-reset, the FCT credit counter shall not be modified.

d. When an FCT is received for a particular virtual channel, 256 shall be added to the FCT credit counter for that virtual channel.

e. When a data frame is sent by a particular virtual channel, the number of N-Chars and Fills sent in the data frame shall be subtracted from the FCT credit counter.

f. A virtual channel shall not be permitted to send data frames when its FCT credit counter is less than 256, unless its output VCB contains an EOP or EEP, in which case it can send N-Chars and Fills up to the limit indicated by the FCT transmit credit counter for that virtual channel.

g. When a data frame is ready to send, the amount of data in the frame shall be compared to the value of the FCT credit counter and if the FCT credit counter is greater than or equal to the amount of data in the frame waiting to be sent, that data frame shall be sent.

h. If the FCT credit counter overflows, the FCT causing the overflows shall be discarded, the FCT credit counter saturated to its maximum value and the error reported in a status register.

5.5.4 Medium access control

a. A medium access controller shall determine which virtual channel is allowed to send the next data frame.

b. Only output virtual channel buffers indicating that they have data ready to form a data frame, shall be permitted to compete for sending the next data frame.

c. Only virtual channels whose FCT credit counter indicates that its input virtual channel buffer at the far end of the SpaceFibre link has room to accept the amount of data to be sent in the next data frame shall be permitted to compete for sending the next data frame.

d. Each output virtual channel buffer, with data ready to send in its output virtual channel buffer and space available in its input virtual channel buffer at the far end of the SpaceFibre link (a ready virtual channel), shall compete for sending the next data frame, if it is allowed to do so in the current time slot.

5.5.4.1 Quality of service

a. Each virtual channel shall support the following qualities of service: best effort, priority, bandwidth reservation, and scheduled.

b. It shall be possible to set the quality of service parameters of each virtual channel individually so that different qualities of service can be applied to different virtual channels.

c. Upon cold reset the quality of service for each virtual channel shall be set to best effort, i.e. lowest priority.

5.5.4.2 Schedule

a. Time shall be split into a number of time-slots of equal duration specified by the time slot duration management parameter, see section 5.14.2.

NOTE Typical time-slot duration would be 0.1 to 1 ms.

b. Time for use in the scheduled quality of service shall be taken from the local time register, which is regularly updated by the time-distribution broadcast channels.

c. The local time register shall provide an indication of when one time-slot finishes and the next one starts, together with the number of that time-slot.

d. The number of time-slots in a schedule shall be specified by a management parameter, see section 5.14.2.

e. The schedule shall indicate in which time-slots a particular virtual channel is permitted to send data frames.

NOTE For each virtual channel there will be a list of time-slots in which it is allowed to send data.

f. It shall be possible for several virtual channels to be scheduled to send data in the same time-slot.

g. If two or more virtual channels are scheduled to send data in the same time-slot, the medium access controller shall send data from the virtual channel with the highest precedence.

h. When a time-slot starts, for which a specific virtual channel is scheduled to send data, has data ready to send in its output virtual channel buffer and has space in its input virtual channel buffer at the far end of the SpaceFibre link, it shall compete with other virtual channels for sending frames over the link based on its precedence.

i. A virtual channel that is not scheduled in the current time-slot shall not be allowed to compete for sending data over the link.

j. At the end of the time-slot, any virtual channel sending data frames shall cease sending them after the current data frame has been sent, unless that particular virtual channel buffer is also scheduled to send data in the next time-slot.

k. The scheduling of traffic from virtual channels over the SpaceFibre link, shall take into account the fact that at the start of a time-slot, there might still be a complete data frame to send over the link from the previous time-slot.

l. When a SpaceFibre network is not required to use the scheduling quality of service, the schedule for each virtual channel shall be set to use all time-slots.

5.5.4.3 Precedence

a. A virtual channel shall compete with other virtual channels for sending frames over the link based on the current precedence of the virtual channel and its schedule.

b. The medium access controller shall use the precedence of each virtual channel to determine which ready virtual channel is permitted to send a data frame.

c. The precedence of each virtual channel shall be compared when the last word of the previous frame is being passed to the framing layer for transmission.

d. The virtual channel buffer that is ready and which has the highest precedence shall be permitted to send the next data frame.

e. The precedence of a virtual channel shall be determined by its quality of service parameters and its bandwidth credit.

f. The precedence of a virtual channel shall be calculated as the sum of the priority precedence plus the current value of bandwidth credit for that virtual channel.

Precedence = Priority Precedence + Bandwidth Credit

NOTE The priority precedence for a virtual channel is derived directly from its priority level parameter.

g. Precedence shall be calculated taking into account the qualities of service mechanisms outlined in Table 5-11.

Table 5-11 Precedence for Different Qualities of Service

Quality of Service	Precedence
Best Effort	A best effort virtual channel is only allowed to send data when no other virtual channel with a different quality of service is ready to send data. If two or more best effort virtual channels are both ready to send data, the one with the highest bandwidth credit is permitted to send a data frame first. Best effort quality of service is obtained when a virtual channel has its priority quality of service parameter set to the lowest priority.
Priority	A priority virtual channel is only allowed to send data when no other virtual channel with a higher priority setting is ready to send data. Several priority levels are supported. If two or more priority virtual channels set to the same priority level are both ready to send data, the one with the highest bandwidth credit is permitted to send a data frame first. If a SpaceFibre network is required to operate using priority only, each virtual channel has to be assigned a different priority level.
Bandwidth Reserved	Bandwidth reservation determines the precedence of a virtual channel based on the link bandwidth reserved for that virtual channel and its recent link utilisation.

	Each virtual channel computes a bandwidth credit based on the link bandwidth reserved for that virtual channel and its recent link utilisation. A virtual channel with large reserved bandwidth and low recent bandwidth utilisation will have high bandwidth credit. A virtual channel with similar reserved bandwidth and higher recent bandwidth utilisation will have lower bandwidth credit. Within a specific priority level virtual channels compete for sending the next data frame based on their current bandwidth credit. A bandwidth reserved virtual channel is only allowed to send data when there are no other virtual channels with higher priority that are ready to send data, and the virtual channel has the highest bandwidth credit of all of the virtual channels of the same priority level that are ready to send data. If a SpaceFibre network that just uses bandwidth reservation is required, this can be achieved by simply setting all virtual channels to use the same priority level.
Scheduled	Scheduled quality of service provides a means of ensuring fully deterministic allocation of SpaceFibre network resources. Time is separated into a series of time-slots during which a virtual channel can be scheduled to send data. When a time-slot arrives in which a virtual channel is scheduled, it can send data based on its precedence. During all the other time-slots when the virtual channel is not scheduled to send data, it is not permitted to send any data even when no other virtual channel has data to send. If a virtual channel scheduled in the current time-slot does not have data to send during its time-slot or does not have space in its VCBs at the far end of the link, other virtual channels allowed to transmit in the current time-slot are permitted to use the otherwise wasted bandwidth. If each virtual channel uses a different slot to send data, a simple scheduled quality of service is provided, without considering priority or bandwidth reservation.

5.5.4.4 Bandwidth credit

a. Bandwidth Credit shall be updated every time a data frame has been sent on any virtual channel.

b. If no data frame is sent, the Bandwidth Credit shall be updated after an interval approximately equal to the time it takes to send a full data frame, or less.

NOTE The clause means that every 66 words or less the Bandwidth Credit is updated.

NOTE When no data is sent the Used Bandwidth is zero, see section 5.5.4.4g.

c. The Available Bandwidth shall be measured for the SpaceFibre links as the number of data or control words that can be sent since the Bandwidth Credit was last updated, allowing for multi-lane operation and taking into account overheads due to framing and control words.

NOTE The overhead is for the frame delimiters is normally 2 words for every 64 data words. Including the FCT and ACK adds another two words of overhead, resulting in then it is 4 every 64 data words which means that the available bandwidth is 64/(64+2+2) = 94% of the link bandwidth. For example a link running at a bit signalling rate of 2.5 Gbits/s the Available Bandwidth is 58.75 Mwords/s.

NOTE When specifying a bandwidth allocation for a virtual channel it is useful to think in terms of the percentage of overall link capacity being allocated. It is then straightforward to allocate a percentage of the bandwidth to each virtual channel, leaving some spare capacity for possible broadcast messages and general contingency.

NOTE When setting the bandwidth allocations each virtual channel is allocated a little bit more bandwidth than it requires, so that the counter is nominally saturated, and the sum of all virtual channel bandwidth allocations is less than 100%, to leave some margin for broadcasts, ACKs, FCT, etc.

d. A virtual channel shall specify a portion of the Link Bandwidth that it wishes or expects to use, i.e. its Expected Bandwidth Portion, including the overhead due to the use of frame delimiters and other control words.

NOTE This information is to be provided for all virtual channels regardless of the quality of service they are providing, using management parameters.

e. The Expected Bandwidth Portion shall be greater than zero for all virtual channels, even when they are not expected to be used.

f. Used Bandwidth shall be the amount of data words sent by a particular virtual channel, including frame delimiters, since the last time the Bandwidth Credit was updated.

NOTE This is zero for all virtual channels except for the one that sent a data frame.

g. Bandwidth Credit shall be calculated independently for each virtual channel as follows:

$$BandwidthCredit = \sum (AvailableBandwidth - \frac{UsedBandwidth}{ExpectedBandwidthPortion})$$

h. Bandwidth Credit shall be permitted to go negative.

NOTE A negative value indicates that the virtual channel is using more than its expected amount of link bandwidth.

i. Bandwidth Credit shall saturate at plus or minus the Bandwidth Credit Limit, i.e. if the Bandwidth Credit reaches a Bandwidth Credit Limit it is set to the value of the Bandwidth Credit Limit.

j. The Bandwidth Credit Limit shall be set to a value not smaller than the Link Bandwidth.

NOTE This leads to a virtual channel to reaching positive saturation after 1 second when no data is being sent by that virtual channel. After this period of time the virtual channel effectively starts to forget what it has send or not sent previously.

k. When the Bandwidth Credit for a virtual channel reaches the positive Bandwidth Credit Limit and stays at that value for at least the Virtual Channel Idle Time Limit, the virtual channel shall indicate in a status register that it is idle and using much less bandwidth than expected.

NOTE A network management application is able to use this information to check correct utilisation of link bandwidth by its various virtual channels.

NOTE The Virtual Channel Idle Time Limit is a configuration parameter that applies to all virtual channels. A typical value is 1 ms.

l. When the Bandwidth Credit for a virtual channel reaches the negative Bandwidth Credit Limit it shall indicate in a status register that the virtual channel is using more bandwidth than expected.

NOTE An implementation might stop a virtual channel from sending more data when the Bandwidth Credit is at the negative Bandwidth Credit Limit, e.g. by disabling the VC transmission.

m. Bandwidth Credit shall be set to zero on cold reset.

n. Bandwidth Credit shall not be altered on warm reset.

o. Upon cold reset the Expected Bandwidth Portion for virtual channel zero (VC0) shall be set to 10% and that of all the other virtual channel set to the minimum possible value.

NOTE After cold reset the local application or a remote application can set the Expected Bandwidth Portion of each virtual channel as required.

NOTE Virtual channel zero (VC0) is used for configuration, control and monitoring of the SpaceFibre network.

5.5.4.5 Priority quality of service

a. There shall be sixteen (TBC) priority levels: 0 to 15, where priority level 0 has the highest precedence and 15 has the lowest precedence.

b. Priority level 15 shall have a Priority Precedence value of P, where P is at least equal to the Bandwidth Credit Limit.

NOTE Precedence = Priority Precedence + Bandwidth Credit

c. Priority level 14 shall have a Priority Precedence value of 3P.

NOTE This means that the precedence at one priority level cannot overlap with that of an adjacent priority level, regardless of the current Bandwidth Credit value.

d. Priority level 13 shall have a Priority Precedence value of 5P.

e. The priority levels shall generally have a Priority Precedence value of 2P(15-N)+P, where N is the priority level, resulting in priority level 1 having Priority Precedence of 29P.

f. Priority level 0 shall have a Priority Precedence value of 31P.

g. Each virtual channel shall be able to be assigned any of the sixteen priority levels.

h. It shall be possible to set more than one virtual channel to the same priority level.

i. When more than one virtual channel has the same priority level, the one with the highest Bandwidth Credit value shall be selected to send the next data frame.

5.5.4.6 Best effort quality of service

a. To configure a virtual channel to provide best effort quality of service the Priority Precedence value shall be set to P.

NOTE This is equivalent to setting the virtual channel to the lowest priority.

5.5.4.7 Bandwidth reservation quality of service

a. Virtual channels, scheduled to send data in the same time-slots, with the same Priority Precedence shall compete based only on the bandwidth credit of that virtual channel.

NOTE If all virtual channels use the same priority level, bandwidth reserved quality of service is used without taking into account any other factor.

5.6 Broadcast message layer

NOTE The Broadcast Message functionality is in the process of being moved to the Network layer.

5.7 Framing layer

5.7.1 Framing layer service interface

The service primitives that shall be associated with this service are:

TX_DATA_FRAME.request;

RX_DATA_FRAME.indication;

TX_FCT.request;

RX_FCT.indication.

TX_BROADCAST_FRAME.request;

RX_BROADCAST_FRAME.indication;

5.7.1.1 TX_DATA_FRAME.request

5.7.1.1.1 Function

The virtual channel layer shall pass a TX_DATA_FRAME.request primitive to the framing layer to request to send a data frame containing some SpaceFibre packet data over the SpaceFibre link.

5.7.1.1.2 Semantics

The TX_DATA_FRAME.request primitive shall provide the following parameters:

TX_DATA_FRAME.request (Virtual Channel, Data)

5.7.1.1.3 When Generated

The TX_DATA_FRAME.request primitive shall be generated when the virtual channel layer has a data to send over the SpaceFibre link.

5.7.1.1.4 Effect On Receipt

On receipt of the TX_DATA_FRAME.request primitive the framing layer shall frame the data from the virtual channel layer forming a data frame and send it over the SpaceFibre link.

5.7.1.2 RX_DATA_FRAME.indication

5.7.1.2.1 Function

The framing layer shall pass a RX_DATA_FRAME.indication primitive to the virtual channel layer to indicate that a data frame has been received.

5.7.1.2.2 Semantics

The RX_DATA_FRAME.indication primitive provides parameters as follows:

RX_DATA_FRAME.indication (Virtual Channel, Data).

5.7.1.2.3 When Generated

The RX_DATA_FRAME.indication primitive shall be passed to the virtual channel layer when a data frame has been received, de-capsulated and unscrambled.

5.7.1.2.4 Effect On Receipt

The effect on receipt of the RX_DATA_FRAME.indication primitive by the virtual channel layer shall be that the SpaceFibre packet data is placed in the virtual channel buffer identified by the virtual channel parameter.

5.7.1.3 TX_FCT.request

5.7.1.3.1 Function

The virtual channel layer shall pass a TX_FCT.request primitive to the framing layer to request to send an FCT over the SpaceFibre link.

5.7.1.3.2 Semantics

The TX_FCT.request primitive shall provide the following parameters:

TX_FCT.request (Virtual Channel).

5.7.1.3.3 When Generated

The TX_FCT.request primitive shall be generated when one of the virtual channel input buffers has room for another data frame.

5.7.1.3.4 Effect On Receipt

On receipt of the TX_FCT.request primitive the framing layer shall immediately send an FCT for the specified virtual channel.

5.7.1.4 RX_FCT.indication

5.7.1.4.1 Function

The framing layer shall pass an RX_FCT.indication primitive to the virtual channel layer to indicate that an FCT has been received.

5.7.1.4.2 Semantics

The RX_FCT.indication primitive provides parameters as follows:

RX_FCT.indication (Virtual Channel).

5.7.1.4.3 When Generated

The RX_FCT.indication primitive shall be passed to the virtual channel layer when an FCT is received.

5.7.1.4.4 Effect On Receipt

The effect on receipt of the RX_FCT.indication primitive by the virtual channel layer shall be to enable another data frame from the specified virtual channel to be sent.

5.7.1.5 TX_BROADCAST_FRAME.request

5.7.1.5.1 Function

The virtual channel layer shall pass a TX_BROADCAST_FRAME.request primitive to the framing layer to request it to frame the broadcast message and to send it over the SpaceFibre link.

5.7.1.5.2 Semantics

The TX_BROADCAST_FRAME.request primitive shall provide the following parameters:

> TX_BROADCAST_FRAME.request (Broadcast Channel, Broadcast Sequence Number, Message).

5.7.1.5.3 When Generated

The TX_BROADCAST_FRAME.request primitive shall be generated when a broadcast message is ready to send.

5.7.1.5.4 Effect On Receipt

On receipt of the TX_BROADCAST_FRAME.request primitive the framing layer shall encapsulate the broadcast message and immediately send it over the SpaceFibre link.

5.7.1.6 RX_BROADCAST_FRAME.indication

5.7.1.6.1 Function

The framing layer shall pass an RX_BROADCAST_FRAME.indication primitive to the broadcast layer to indicate that a broadcast frame has arrived and to pass its contents to the broadcast layer.

5.7.1.6.2 Semantics

The RX_BROADCAST_FRAME.indication primitive provides parameters as follows:

> RX_BROADCAST_FRAME.indication (Broadcast Channel, Broadcast Sequence Number, Message).

5.7.1.6.3 When Generated

The RX_BROADCAST_FRAME.indication primitive shall be passed to the broadcast layer when a broadcast frame is received.

5.7.1.6.4 Effect On Receipt

The effect on receipt of the RX_BROADCAST_FRAME.indication primitive by the broadcast layer shall be for the broadcast layer to validate the broadcast frame and to pass valid broadcast messages up to the user application.

5.7.2 Framing

a. The framing layer shall encapsulate data from the virtual channel layer into data frames, scramble, and pass them to the retry layer.

b. The framing layer shall receive flow control information from the virtual channel layer and pass it to the retry layer.

c. The framing layer shall encapsulate data from the broadcast message layer into broadcast frames and pass them to the retry layer.

d. Data frames from the retry layer shall be de-capsulated, unscrambled and passed up to the virtual channel layer.

e. FCTs received from the retry layer shall be passed up to the virtual channel layer.

f. Broadcast frames from the retry layer shall be de-capsulated and passed up to the broadcast channel layer.

5.7.3 EM emission mitigation

5.7.3.1 Data scrambling

a. The data field of data frames shall be scrambled prior to transmission of the frame by bit-wise multiplication of the data with a sequence of random numbers produced from a scrambling polynomial.

NOTE Bit-wise multiplication is the XOR function.

NOTE This scrambler is illustrated in Figure 5-5, which is included for clarity. The scrambler can be implemented in other ways.

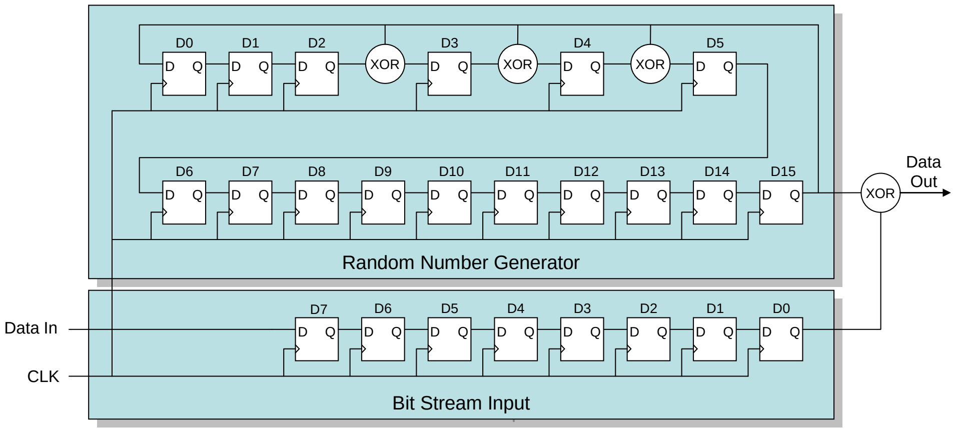

Figure 5-5 Scrambler / De-Scrambler

b. The scrambling polynomial to use shall be $G(x) = X^{16} + X^5 + X^4 + X^3 + 1$

c. The seed for the scrambler shall be 0xffff i.e. all flip-flops in the random number generator are set to one.

d. The scrambler shall be re-seeded at the start of every new data frame.

e. The single bit output sequence from the random number generator shall be XORed with the bit sequence from the data word.

f. The least-significant bit of the least-significant character in each word shall be scrambled first.

g. When an EOP or EEP or Fill occurs in the data it shall not be replaced by scrambled data but the scrambler shall continue to generate random numbers normally.

NOTE If the EOP or EEP or Fill control code was replaced by scrambled data it would not be possible to distinguish it from a data value.

5.7.3.2 Unscrambling

a. When data frames are received they shall be unscrambled by multiplying (XORing) the received data with the same scrambling sequence as was used for scrambling.

b. The control codes representing EOP or EEP shall not be unscrambled.

5.7.4 Frame reception

a. Data frames that are received without error and in the correct order shall be passed from the retry layer to the framing layer.

b. FCTs received without error and in the correct order shall be passed from the retry layer to the framing layer.

c. Broadcast frames that are received without error and in the correct order shall be passed from the retry layer to the framing layer.

5.8 Retry layer

5.8.1 Retry layer service interface

The service primitives that shall be associated with the retry layer service are:

TX_DATA_FRAME.request;
TX_DATA_RDY.indication;
RX_DATA_FRAME.indication;
TX_FCT.request;
TX_FCT_RDY.indication;
RX_FCT.indication;
TX_BC_FRAME.request;
TX_BC_RDY.indication;
RX_BC_FRAME.indication;

5.8.1.1 TX_DATA_FRAME.request

5.8.1.1.1 Function

The framing layer shall pass a TX_DATA_FRAME.request primitive to the retry layer to send a data frame over the SpaceFibre link.

5.8.1.1.2 Semantics

The TX_DATA_FRAME.request primitive shall provide the following parameters:

TX_DATA_FRAME.request (Data Frame)

5.8.1.1.3 When Generated

The TX_DATA_FRAME.request primitive shall be passed to the retry layer when the framing layer has a data frame to send over the SpaceFibre link.

5.8.1.1.4 Effect On Receipt

On receipt of the TX_DATA_FRAME.request primitive the retry layer shall send the data frame over the SpaceFibre interface, resending the data frame in the event of it failing to be received correctly at the far end of the SpaceFibre link.

5.8.1.2 TX_DATA_RDY.indication

5.8.1.2.1 Function

The retry layer shall pass a TX_DATA_RDY.indication primitive to the framing layer to indicate that it is ready to accept a new data frame.

5.8.1.2.2 Semantics

The TX_DATA_RDY.indication primitive shall have no parameters.

5.8.1.2.3 When Generated

The TX_DATA_RDY.indication primitive shall be passed to the framing layer when the retry layer can accept a new request to transfer a data frame.

5.8.1.2.4 Effect On Receipt

On receipt of the TX_DATA_RDY.indication primitive the framing layer shall be permitted to send a new TX_DATA_FRAME.request when it is ready to do so.

5.8.1.3 RX_DATA_FRAME.indication

5.8.1.3.1 Function

The retry layer shall pass a RX_DATA_FRAME.indication primitive to the framing layer to indicate that a data frame has been received.

5.8.1.3.2 Semantics

The RX_DATA_FRAME.indication primitive provides parameters as follows:

RX_DATA_FRAME.indication (Data Frame).

5.8.1.3.3 When Generated

The RX_DATA_FRAME.indication primitive shall be passed to the framing layer when a data frame is received without error and in the correct order.

5.8.1.3.4 Effect On Receipt

The effect on receipt of the RX_DATA_FRAME.indication primitive by the framing layer shall be that the contents of the received data frame is extracted and passed to the virtual channel layer.

5.8.1.4 TX_FCT.request

5.8.1.4.1 Function

The framing layer shall pass a TX_FCT.request primitive to the retry layer to send an FCT over the SpaceFibre link.

5.8.1.4.2 Semantics

The TX_FCT.request primitive shall provide the following parameters:

TX_FCT.request (FCT)

5.8.1.4.3 When Generated

The TX_FCT.request primitive shall be passed to the retry layer when the framing layer has a FCT to send over the SpaceFibre link.

5.8.1.4.4 Effect On Receipt

On receipt of the TX_FCT.request primitive the retry layer shall send the FCT over the SpaceFibre interface, resending the FCT in the event of it failing to be received correctly at the far end of the SpaceFibre link.

5.8.1.5 TX_FCT_RDY.indication

5.8.1.5.1 Function

The retry layer shall pass a TX_FCT_RDY.indication primitive to the framing layer to indicate that it is ready to accept a new FCT.

5.8.1.5.2 Semantics

The TX_FCT_RDY.indication primitive shall have no parameters.

5.8.1.5.3 When Generated

The TX_FCT_RDY.indication primitive shall be passed to the framing layer when the retry layer can accept a new request to transfer an FCT.

5.8.1.5.4 Effect On Receipt

On receipt of the TX_FCT_RDY.indication primitive the framing layer shall be permitted to send a new TX_FCT_FRAME.request when it is ready to do so.

5.8.1.6 RX_FCT.indication

5.8.1.6.1 Function

The retry layer shall pass an RX_FCT.indication primitive to the framing layer to indicate that an FCT has been received.

5.8.1.6.2 Semantics

The RX_FCT.indication primitive provides parameters as follows:

RX_FCT.indication (FCT).

5.8.1.6.3 When Generated

The RX_FCT.indication primitive shall be passed to the framing layer when an FCT is received without error and in the correct order.

5.8.1.6.4 Effect On Receipt

The effect on receipt of the RX_FCT.indication primitive by the framing layer shall be for the information contained in the FCT to be extracted and passed on up to the virtual channel layer.

5.8.1.7 TX_BC_FRAME.request

5.8.1.7.1 Function

The framing layer shall pass a TX_BC_FRAME.request primitive to the retry layer to send a broadcast frame over the SpaceFibre link.

5.8.1.7.2 Semantics

The TX_BC_FRAME.request primitive shall provide the following parameters:

TX_BC_FRAME.request (Broadcast Frame)

5.8.1.7.3 When Generated

The TX_BC_FRAME.request primitive shall be passed to the retry layer when the framing layer has a broadcast frame to send over the SpaceFibre link.

5.8.1.7.4 Effect On Receipt

On receipt of the TX_BC_FRAME.request primitive the retry layer shall immediately send the broadcast frame over the SpaceFibre link, resending the broadcast frame in the event of it failing to be received correctly at the far end of the SpaceFibre link.

5.8.1.8 TX_BC_RDY.indication

5.8.1.8.1 Function

The retry layer shall pass a TX_BC_RDY.indication primitive to the framing layer to indicate that it is ready to accept a new broadcast frame.

5.8.1.8.2 Semantics

The TX_BC_RDY.indication primitive shall have no parameters.

5.8.1.8.3 When Generated

The TX_BC_RDY.indication primitive shall be passed to the framing layer when the retry layer can accept a new request to transfer a broadcast frame.

5.8.1.8.4 Effect On Receipt

On receipt of the TX_BC_RDY.indication primitive the framing layer shall be permitted to send a new TX_BC_FRAME.request when it is ready to do so.

5.8.1.9 RX_BC_FRAME.indication

5.8.1.9.1 Function

The retry layer shall pass a RX_BC_FRAME.indication primitive to the framing layer to indicate that a broadcast frame has been received.

5.8.1.9.2 Semantics

The RX_BC_FRAME.indication primitive provides parameters as follows:

RX_BC_FRAME.indication (Broadcast Frame).

5.8.1.9.3 When Generated

The RX_BC_FRAME.indication primitive shall be passed to the framing layer when a broadcast frame is received without error and in the correct order.

5.8.1.9.4 Effect On Receipt

The effect on receipt of the RX_BC_FRAME.indication primitive by the framing layer shall be that the contents of the received broadcast frame is extracted and passed to the broadcast layer.

5.8.2 Data words identification state machine

a. The Data Word Identification state machine shall be used to identify which type of frame a data word belongs to, in order to support de-multiplexing of different frame types.

NOTE The state diagram for the Data Word Identification state machine is illustrated in Figure 5-6.

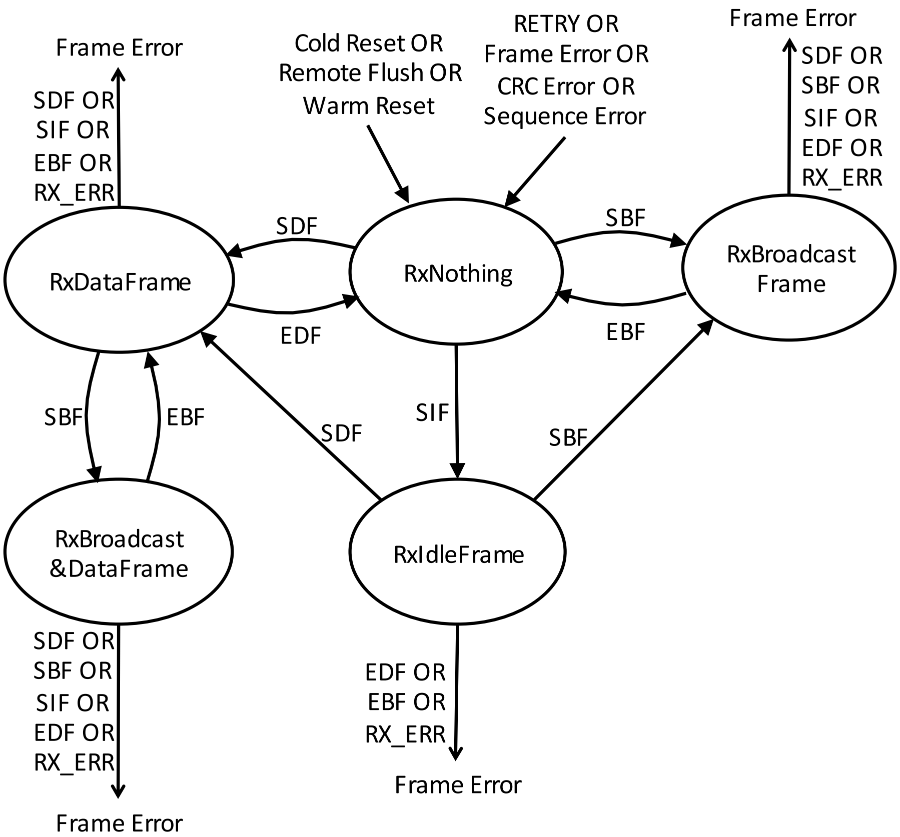

Figure 5-6 Data Word Identification State Machine

5.8.2.2 RxNothing state

a. The RxNothing state shall be entered on one of the following conditions:

1. Following a cold reset, warm reset or remote flush indication.
2. From the RxDataFrame state, when an EDF is received.
3. From the RxBroadcastFrame state, when an EBF is received.
4. From any state when a RX_ERROR, a RETRY, a sequence error, a CRC error or an unexpected control word (Frame Error) is received.

NOTE A sequence error occurs when the frame sequence number of a data frame, FCT or broadcast frame is not one more (using modulo 128) than the value in the receive frame sequence counter. It also occurs when the frame sequence number of an idle frame does not match the value of the receive frame sequence counter.

b. When in the RxNothing state the Data Word Identification state machine shall initiate the following actions:

1. Indicate that no frame data is being received.
2. Discard any EBF and EDF that are received.

c. The Data Word Identification state machine shall leave the RxNothing State on one of the following conditions, which shall be evaluated in the order given:

1. When an SBF is received, move to the RxBroadcastFrame state.
2. When an SDF is received, move to the RxDataFrame state.
3. When an SIF is received, move to the RxIdleFrame state.

NOTE No action is taken in the RxNothing state when an EDF or EBF is received.

d. The RxNothing state is summarised in Table 5-12.

Table 5-12 RxNothing State

State	RxNothing
Entry	Following a cold reset or warm reset or remote flush From the RxDataFrame state, when an EDF is received. From the RxBroadcastFrame state, when an EBF is received. From any state when a RX_ERROR, a RETRY, a sequence error a CRC error or an unexpected control word (Frame Error) is received.
Action	Indicate that no frame data is being received.
Exit	When an SBF is received, move to the RxBroadcastFrame state. When an SDF is received, move to the RxDataFrame state. When an SIF is received, move to the RxIdleFrame state.

5.8.2.3 RxDataFrame state

a. The RxDataFrame state shall be entered on one of the following conditions:

1. From the RxNothing state, when an SDF is received.
2. From the RxIdleFrame state, when an SDF is received.
3. From the RxBroadcast&DataFrame state, when an EBF is received.

b. When in the RxDataFrame state the Data Word Identification state machine shall initiate the following actions:

1. Indicate that data words being received belong to a data frame.

c. The Data Word Identification state machine shall leave the RxDataFrame state on one of the following conditions, which shall be evaluated in the order given:

1. When an RETRY is received, move to the RxNothing state.
2. When an RXERR is received, move to the RxNothing state.
3. When a CRC error is detected, move to the RxNothing state.
4. When a sequence error is detected, move to the RxNothing state.
5. When an SBF is received move to the RxBroadcast&DataFrame state.
6. When an EDF is received, move to the RxNothing state.
7. When an SDF, SIF or EBF is received (Frame Error) move to the RxNothing state.

d. The RxDataFrame state is summarised in Table 5-13.

Table 5-13 RxDataFrame State

State	RxDataFrame
Entry	From the RxNothing state, when an SDF is received. From the RxIdleFrame state, when an SDF is received. From the RxBroadcast&DataFrame state, when an EBF is received.
Action	Indicate that data words being received belong to a data frame.
Exit	When an RETRY is received, move to the RxNothing state. When an RXERR is received, move to the RxNothing state. When a CRC error is detected, move to the RxNothing state. When a sequence error is detected, move to the RxNothing state. When an SBF is received move, to the RxBroadcast&DataFrame state. When an EDF is received, move to the RxNothing state. When an SDF, SIF or EBF is received (Frame Error) move to the RxNothing state.

5.8.2.4 RxBroadcastFrame state

a. The RxBroadcastFrame state shall be entered on one of the following conditions:

1. From the RxNothing state, when an SBF is received.
2. From the RxIdleFrame state, when an SBF is received.

b. When in the RxBroadcastFrame state the Data Word Identification state machine shall initiate the following actions:

1. Indicate that data words being received belong to a broadcast frame.
2. Count the number of data words being received.

c. The Data Word Identification state machine shall leave the RxBroadcastFrame state on one of the following conditions, which shall be evaluated in the order given:

1. When an RETRY is received, move to the RxNothing state.
2. When an RXERR is received, move to the RxNothing state.
3. When a CRC error is detected, move to the RxNothing state.
4. When a sequence error is detected, move to the RxNothing state.
5. When an EBF is received, move to the RxNothing state.
6. When an SDF, SBF, SIF or EDF is received (Frame Error) move to the RxNothing state.

d. The RxBroadcastFrame state is summarised in Table 5-14.

Table 5-14 RxBroadcastFrame State

State	RxBroadcastFrame
Entry	From the RxNothing state, when an SBF is received. From the RxIdleFrame state, when an SBF is received.
Action	Indicate that data words being received belong to a broadcast frame. Count the number of data words being received.
Exit	When an RETRY is received, move to the RxNothing state. When an RXERR is received, move to the RxNothing state. When a CRC error is detected, move to the RxNothing state. When a sequence error is detected, move to the RxNothing state. When an EBF is received, move to the RxNothing state. When an SDF, SBF, SIF or EDF is received (Frame Error) move to the RxNothing state.

5.8.2.5 RxBroadcast&DataFrame state

a. The RxBroadcast&DataFrame state shall be entered on the following condition:

1. From the RxDataFrame state, when an SBF is received.

b. When in the RxNothing state the Data Word Identification state machine shall initiate the following actions:

1. Indicate that data words being received belong to a broadcast frame.
2. Count the number of data words being received.

c. The Data Word Identification state machine shall leave the RxBroadcast&DataFrame state on one of the following conditions, which shall be evaluated in the order give:

1. When an RETRY is received, move to the RxNothing state.
2. When an RXERR is received, move to the RxNothing state.
3. When a CRC error is detected, move to the RxNothing state.
4. When a sequence error is detected, move to the RxNothing state.
5. When an EBF is received and the data word counter is equal to two, move to the RxDataFrame state.
6. When an EBF is received and the data word counter is not equal to two, move to the RxNothing state.
7. When an SDF, SBF, SIF or EDF is received (Frame Error) move to the RxNothing state.

d. The RxBroadcast&DataFrame state is summarised in Table 5-15.

Table 5-15 RxBroadcast&DataFrame State

State	RxBroadcast&DataFrame
Entry	From the RxDataFrame state, when an SBF is received.
Action	Indicate that data words being received belong to a broadcast frame.
Exit	When an RETRY is received, move to the RxNothing state. When an RXERR is received, move to the RxNothing state. When a CRC error is detected, move to the RxNothing state. When a sequence error is detected, move to the RxNothing state. When an EBF is received and the data word counter is equal to two, move to the RxDataFrame state. When an EBF is received and the data word counter is not equal to two, move to the RxNothing state. When an SDF, SBF, SIF or EDF is received (Frame Error) move to the RxNothing state.

5.8.2.6 RxIdleFrame state

a. The RxIdleFrame state shall be entered on one of the following conditions:

1. From the RxNothing state, when an SIF is received.

b. When in the RxIdleFrame state the Data Word Identification state machine shall initiate the following actions:

1. Indicate that data words being received belong to an idle frame.

c. The Data Word Identification state machine shall leave the RxIdleFrame state on one of the following conditions, which shall be evaluated in the order give:

1. When an RETRY is received, move to the RxNothing state.
2. When an RXERR is received, move to the RxNothing state.
3. When an EDF or EBF is received (Frame Error), move to the RxNothing state.
4. When an SBF is received, move to the RxBroadcastFrame state.
5. When an SDF is received, move to the RxDataFrame state.
6. When an SIF is received, remain in the RxIdleFrame state.

d. The RxIdleFrame state is summarised in Table 5-16.

Table 5-16 RxIdleFrame State

State	RxIdleFrame
Entry	From the RxNothing state, when an SIF is received.
Action	Indicate that data words being received belong to an idle frame.
Exit	When an RETRY is received, move to the RxNothing state. When an RXERR is received, move to the RxNothing state. When an EDF or EBF is received (Frame Error), move to the RxNothing state. When an SBF is received, move to the RxBroadcastFrame state. When an SDF is received, move to the RxDataFrame state. When an SIF is received, remain in the RxIdleFrame state.

5.8.2.7 Control Words

a. It shall be possible to receive FCT, ACK, NACK, FULL, and RETRY control words in any state.

5.8.3 Receive frame buffering

a. While a data frame or broadcast frame is being received, their words shall be buffered in order to be able to check the CRC and frame sequence number prior to the frame being accepted and passed up to the Framing layer.

b. Only data frames, FCTs, and broadcast frames that are received without error and in the correct sequence shall be passed up to the framing layer.

c. Data frames, FCTs, and broadcast frames that are received containing errors or in an incorrect sequence shall be discarded.

d. A data frame receive buffer shall store data words received in the RxDataFrame state of the Data Word Identification state machine.

e. A broadcast frame receive buffer shall store data words received in the RxBroadcastFrame and RxBroadcast&DataFrame state of the Data Word Identification state machine.

f. The data frame receive buffer shall be flushed when a SDF control word is received.

g. The broadcast frame receive buffer shall be flushed when a SBF control word is received.

h. The data and broadcast frame receive buffers shall be flushed when the Data Word Identification state machine enters the RxNothing state.

i. A data frame shall be passed up to the framing layer for de-framing when an EDF is received without errors.

j. A broadcast frame shall be passed up to the framing layer for de-framing when an EBF is received without errors and the broadcast frame contains two data words.

k. When a cold reset, warm reset or remote flush is requested the data frame receive buffer and broadcast frame receive buffer, shall be emptied.

5.8.4 CRC for data frame

a. A 16-bit CRC checksum shall be applied to data frames.

b. The CRC shall cover an entire data frame, from and including the comma in the start of data frame control word, up to and including the frame sequence number in the end of data frame control word.

c. When a control code is to be included in the CRC calculation, its data value shall be used.

NOTE For example, K28.7 the comma code is replaced by D28.7 in the CRC calculation.

d. When an EOP or EEP or Fill occurs in the data the value of the control code representing the EOP or EEP or Fill shall be included in the CRC generation.

NOTE For example, an EOP is represented by K29.7 so the value D29.7 would be included in the CRC calculation. Including an EOP or EEP or Fill in the CRC calculation means that they are protected by the CRC checksum.

e. The 16-bit CRC calculation shall be the CRC-16-CCITT.

f. The least-significant bit of the least-significant character in each word shall be scrambled first.

5.8.5 CRC for broadcast frame, FCT, ACK and NACK

a. 8-bit CRC An 8-bit CRC checksum shall be applied to broadcast frames, FCTs, FULLs, ACKs and NACKs.

b. The 8-bit CRC shall cover an entire broadcast frame, from and including the comma in the start of broadcast frame control word, up to and including the frame sequence number in the end of broadcast frame control word.

c. When a control code is to be included in the CRC calculation, its data value shall be used.

NOTE For example, K28.7 the comma code is replaced by D28.7 in the CRC calculation.

d. The 8-bit CRC calculation procedures shall:

1. use modulo 2 arithmetic for polynomial coefficients;
2. use a systematic binary *(n+16,n)* block code, where *(n+8)* is the number of bits of the codeword $c(x)$ and *n* is divisible by 8; *n* is the number of bits covered by the CRC;
3. use the following generating polynomial:

$$g(x) = x^8 + x^2 + x + 1$$

4. use byte format as input and output, for which the bits are represented as:

$$b_7\, b_6\, b_5\, b_4\, b_3\, b_2\, b_1\, b_0$$

where b_7 is the most significant bit and b_0 is the least significant bit;

e. The 8-bit CRC generation procedure shall behave as follows:

1. The procedure accepts an *n*-bit input which is used to construct $m(x)$, where:

(a) the *n*-bit input is defined to be the set of bits $B_{i,j}$ grouped into $n/8$ bytes where $i = \{0,1,\ldots,n/8-1\}$ is the byte index and $j = \{7,6,\ldots,0\}$ is the bit index;

(b) the $n/8$ input bytes correspond to the fields covered by the CRC excluding the CRC byte; the first byte transmitted has index $i = 0$; the last byte transmitted has index $i = n/8-1$;

(c) $m(x)$ is a polynomial $m_{n-1}x^{n-1} + m_{n-2}x^{n-2} + ... + m_0x^0$ having binary coefficients m_i;

(d) $m(x)$ can be represented as an *n*-bit vector where coefficient m_{n-1} of the highest power of x is the most significant bit and coefficient m_0 of the lowest power of x is the least significant bit;

(e) the bit vector representation of $m(x)$ is formed by concatenating the $n/8$ bytes of the input in transmission order, where the least significant bit b_0 of each byte is taken first and the most significant bit b_7 of each byte is taken last:

$$m_{n-1} = B_{0,0}, m_{n-2} = B_{0,1}, m_{n-3} = B_{0,2}, \ldots, m_{n-7} = B_{0,6}, m_{n-8} = B_{0,7},$$
$$m_{n-9} = B_{1,0}, m_{n-10} = B_{1,1}, m_{n-11} = B_{1,2}, \ldots, m_{n-15} = B_{1,6}, m_{n-16} = B_{1,7}, \ldots$$
$$m_7 = B_{n/8-1,0}, m_6 = B_{n/8-1,1}, m_5 = B_{n/8-1,2}, \ldots, m_1 = B_{n/8-1,6}, m_0 = B_{n/8-1,7}$$

2. The procedure generates the remainder polynomial $r(x)$ given by the equation:

$$r(x) = \left[m(x) \cdot x^8 \right] \text{modulo } g(x)$$

where $r(x) = r_7 x^7 + r_6 x^6 + \ldots r_0 x^0$ and r_i are binary coefficients;

3. The Header and Data CRC are formed from the 8-bit vector representation of $r(x)$; the least significant bit b_0 of the CRC byte is coefficient r_7 of the highest power of x, while the most significant bit b_7 of the CRC byte is coefficient r_0 of the lowest power of x:

$$b_7 = r_0, b_6 = r_1, b_5 = r_2, b_4 = r_3, b_3 = r_4, b_2 = r_5, b_1 = r_6, b_0 = r_7$$

NOTE 1 The codeword $c(x) = m(x) \cdot x^8 + r(x)$ is formed by concatenating the bit vector representations of $m(x)$ and $r(x)$.

NOTE 2 When a Galois version of a Linear Feedback Shift Register is used for CRC generation, its initial value is zero.

NOTE 3 Example VHDL and C-code implementations of this CRC algorithm are included in clause Annex B.

f. If the CRC generation procedure is applied to the bytes covered by the CRC *excluding* the CRC byte then the generated CRC shall be compared directly with the expected CRC byte. If the generated and expected CRC bytes are equal then no errors have been detected; if they are different then an error has been detected.

g. If the CRC generation procedure is applied to the bytes covered by the CRC *including* the CRC byte then the output of the CRC generation procedure shall be zero if no errors have been detected and non-zero if an error has been detected.

NOTE 1 When the codeword $c^*(x)$ is input to the CRC generator then the remainder is the syndrome: $s(x) = \left[c^*(x) \cdot x^8 \right] \text{modulo } g(x)$.

NOTE 2 The codeword $c^*(x)$ is the concatenation of the Header or Data bytes covered by the CRC, followed by the CRC byte.

5.8.6 Frame sequence numbering

5.8.6.1 Frame sequence numbers on transmit

a. A single 7-bit transmit frame sequence counter shall be maintained in the transmit side of the SpaceFibre CODEC to hold the 7-bit frame sequence number of the last data frame, broadcast frame or FCT control word that was sent.

b. The 7-bit transmit frame sequence counter shall operate modulo 128.

c. A transmit frame polarity flag shall be associated with the 7-bit transmit frame sequence counter together forming the frame sequence number (FR_SEQ#), see section 5.3.1.4.2.

d. The transmit frame sequence counter shall be cleared to zero on cold reset of the SpaceFibre CODEC.

e. The transmit frame polarity flag shall be cleared to zero on cold reset of the SpaceFibre CODEC.

f. The transmit frame sequence counter shall be cleared to zero on remote flush indication.

g. The transmit frame polarity flag shall be cleared to zero on remote flush indication.

h. The transmit frame sequence counter shall not be modified on warm reset of the SpaceFibre CODEC.

i. The transmit frame polarity flag shall not be modified on warm reset of the SpaceFibre CODEC.

j. The transmit frame polarity flag shall be inverted each time a new retry operation begins.

k. The transmit frame sequence counter and associated polarity flag shall provide frame sequence numbers for data frames, FCTs, broadcast frames, idle frames and FULL words.

l. Immediately prior to an end of data frame, end of broadcast frame, or FCT control word being passed from the retry layer to the lane layer, the current value of the transmit frame sequence counter shall be incremented and the new value together with the current value of the polarity flag placed in the frame sequence number (FR_SEQ#) field of the end of data frame, end of broadcast frame, or FCT control word.

NOTE Every 128 sequence numbers the counter will roll over repeating the series of sequence numbers.

m. When an idle frame is passed from the retry layer to the lane layer, the current value of the transmit frame sequence counter together with the current value of the polarity flag shall be placed in the frame sequence number (FR_SEQ#) field of the start of idle frame (SIF) control word.

5.8.6.2 Frame sequence numbers on reception

a. A single 7-bit receive frame sequence counter shall be maintained in the receive side of the SpaceFibre CODEC to hold the 7-bit frame sequence number of the last data frame, broadcast frame or FCT control word that was received correctly.

b. The 7-bit receive frame sequence counter shall operate modulo 128.

c. A receive frame polarity flag shall be associated with the 7-bit receive frame sequence counter together called the Local Receive Counter, the value of which is checked against received frame sequence numbers (FR_SEQ#).

d. The receive frame sequence counter shall be cleared to zero on cold reset of the SpaceFibre CODEC.

e. The receive frame polarity flag shall be cleared to zero on cold reset of the SpaceFibre CODEC.

f. The receive frame sequence counter shall be cleared to zero on remote flush indication.

g. The receive frame polarity flag shall be cleared to zero on remote flush indication.

h. The receive frame sequence counter shall not be modified on warm reset of the SpaceFibre CODEC.

i. The receive frame polarity flag shall not be modified on warm reset of the SpaceFibre CODEC.

j. The 7-bit receive frame sequence counter together with the receive frame polarity flag shall be used to check the sequence numbers of data frames, FCTs broadcast frames, idle frames and FULL control words.

k. When an EDF, EBF, SIF, FCT, or FULL control word with a valid CRC passes into the retry layer from the multi-lane layer, its frame sequence number (FR_SEQ#) shall be checked against the current value of the receive frame sequence counter and receive frame polarity flag.

l. If the frame sequence number (FR_SEQ#) of a received EDF, EBF or FCT is one more than the local receive counter and the polarity flag of the FR_SEQ# matches the receive frame polarity flag, the data frame, broadcast frame or FCT shall be accepted.

NOTE A sequence error occurs when the frame sequence number of a data frame, broadcast frame or FCT is not one more (using modulo 128) than the value in the receive frame sequence counter or the polarity flag is different. It also occurs when the frame sequence number of an idle frame or FULL control word does not match the value of the receive frame sequence counter.

m. If the frame sequence number (FR_SEQ#) of a received SIF or FULL has the same value as the receive frame sequence counter and their polarity flags match, the idle frame shall be accepted.

NOTE An idle or FULL frame carries the frame sequence number of the last data frame, broadcast frame or FCT transmitted.

n. When a data frame, FCT or broadcast frame is accepted, the receive frame sequence counter shall be incremented.

o. If the frame sequence number (FR_SEQ#) of a received EDF, EBF or FCT is not one more than the receive frame sequence counter, the data frame, FCT or broadcast frame shall be rejected.

p. When a data frame, FCT or broadcast frame is rejected, the receive frame sequence counter shall not be incremented.

q. If the polarity flags of the received FR_SEQ# and the local receive counter do not match, the data frame, FCT, broadcast or idle frame shall be rejected.

r. The polarity flag shall be inverted when either or both of the following conditions occur:

1. The condition specified in the Retry layer error state machine (section 5.7.9.3.1 b) occurs;
2. The CRC is valid, a sequence error occurs, and the polarity flags of the received FR_SEQ# and the local receive counter match.

NOTE This covers the case of frames resent being lost after a retry operation. Inverting the polarity flag will produce a new retry operation.

5.8.7 Frame retry buffering

a. Data frames, FCTs and broadcast frames shall be placed in dedicated frame retry buffers, as they pass from the frame layer to the retry layer, and are sent over the SpaceFibre link.

NOTE The retry buffers are not intermediate buffers: they only contain frames/FCTs that have been sent over the SpaceFibre link.

NOTE The total number of frames or FCT words that can be stored in the frame retry buffers after they are sent is less than 128, because of the size of FR_SEQ#.

b. The data frames, FCTs, and broadcast frames in the retry buffers shall be referenced by the frame sequence number (FR_SEQ#) given to them when they are sent over the SpaceFibre link.

c. When an ACK is received, all data frames, FCTs and broadcast frames, sent with a frame sequence number less than or equal to the frame sequence number in the ACK shall be deleted from the frame retry buffers, taking into account counter roll over effects (i.e. calculation is modulo 128).

d. A NACK received with the polarity flag in the frame sequence number (FR_SEQ#) which is different to that in the transmit frame sequence counter shall be ignored and not cause a retry.

NOTE This prevents an error from causing multiple retries.

e. When a NACK is received with a valid CRC and with a polarity flag in the frame sequence number field that has the same value as the polarity flag of the local transmit frame sequence counter, the NACK shall be accepted.

f. When a NACK is accepted, all the data frames, FCTs, and broadcast frames with a sequence number less than or equal to that of the NACK shall be removed from the frame retry buffer, since the NACK has signalled that they have been received correctly.

g. When a NACK is accepted, the data frame, FCT or broadcast frame, which has a sequence number one more than that of the NACK shall be resent together with all data frames, FCTs and broadcast frames which have higher value sequence numbers, since the NACK has signalled that they have not been received correctly.

NOTE See section 5.8.9.2 for reasons why a NACK is transmitted.

h. When the contents of the retry buffers are resent, the broadcast frames should always be sent before FCT and data frames, even if the broadcast frames were not sent before the NACK was received. The frame precedence detailed in section 5.3.4 shall prevail.

i. When the contents of the retry buffers are resent they shall be given new frame sequence numbers following on from the last correctly received frame sequence number plus one and incrementing with each frame/FCT that is retried or otherwise sent.

j. The data frames, FCTs, and broadcast frames that are resent shall have the polarity flag in the FR-SEQ# set to the new polarity value of the local frame sequence counter.

k. The number of retries shall be recorded and made available in a status register.

l. On cold reset and remote flush the frame/FCT retry buffers shall be emptied.

m. On a cold reset and remote flush the number of retries shall be cleared to zero.

n. On warm reset the frame/FCT retry buffers shall not be emptied

o. On a warm reset the number of retries shall not be cleared.

p. The size of the frame retry buffers shall be implementation dependent.

NOTE The frame retry buffer should be long enough to hold all the data and control words that can fit on the line for as long as it takes to send a frame and receive an ACK in reply.

NOTE When multiple lanes are used proportionally more data can be held on the line.

q. When a retry buffer becomes full, no more data for this buffer shall be accepted from the framing layer.

r. When one of the retry buffers becomes full, FULL control words shall be sent following the precedence detailed in section 5.3.4.

s. FULL words shall only be sent if one or more of the retry buffers are full.

t. When all the frame retry buffers are full, only FULLs, ACKs and NACKs shall be sent by the retry layer.

u. If a broadcast frame is resent during a retry the LATE flag in the RSVD/LATE field of the end of broadcast frame (EBF) shall be set to one, otherwise the LATE flag shall be zero.

NOTE The LATE flag set to one indicates that the broadcast frame has been delayed.

5.8.8 Idle frames

a. Idle frames shall be generated when there are no data frames, FCTs, broadcast frames or FULLs to send.

b. The frame sequence number field in the start of idle frame shall be set to the current value of the local transmit frame sequence counter, including the polarity flag.

NOTE After a cold reset or remote flush this counter is set to zero so if an idle frame is the first frame sent after cold reset, its frame sequence number field will be zero. If a data frame, FCT, or broadcast frame is the first frame sent after a cold reset its frame sequence number field will be one.

c. The start of idle frame control word shall be followed by a series of pseudo-random data words which will form a pseudo random bit sequence when transmitted.

d. The algorithm used to generate the series of pseudo-random data words for the idle frames shall be the same as is used to perform data scrambling, see clause 5.7.3.1.

e. At the end of an idle frame the next pseudo-random data word shall be generated and stored for use as the seed for the next idle frame.

f. When the next idle frame is to be sent the stored seed value shall be used as the first pseudo-random data word, i.e. the seed, in that idle frame.

g. After cold reset or remote flush the stored seed shall be set to 0xffff ffff (TBC).

h. After warm reset the stored seed shall not be modified.

5.8.9 ACK/NACK control

5.8.9.1 Sending ACKs

a. When a data frame, broadcast frame, FCT or FULL is received without a sequence error or CRC error, an ACK shall be requested to be sent

b. ACKs shall be sent as soon as possible after being requested, taking into account control word precedence, see section 5.3.4.

NOTE ACKs are not buffered before they are sent

c. After an ACK is sent there shall be a minimum of 15 words before another ACK is sent.

NOTE This ensures that the bandwidth reserved quality of service works independently of the size of data frames.

d. If another data frame, broadcast frame, FCT or FULL is received while an ACK is pending, only one ACK shall be sent.

NOTE An ACK need not be sent for every data frame, broadcast frame, FCT or FULL.

e. When an ACK is sent, it shall contain the current value of the receive frame sequence counter and the receive frame polarity flag.

f. When a NACK is requested to be sent, any pending ACK shall be cancelled and not sent.

5.8.9.2 Sending NACKs

a. A NACK shall be requested to be sent under any of the following conditions:

1. When a RX_ERR control word is received, or a CRC error or a sequence error occurs while the Data Word Initialisation state machine is in the RxDataFrame, RxBroadcstFrame, or RxData&BroadcastFrame states,

2. When the type of frame received cannot be identified or if an unrecognised word is received or if a word is received in an unexpected place in a frame, or

3. When a received idle frame or a FULL control word contains a frame sequence number which is not the same as the receive frame sequence counter.

NOTE A NACK requested to be sent means that a NACK will be pending to be sent if it cannot be sent immediately.

NOTE The idle frame sequence number ought to have the same sequence number as the last correctly received data frame, FCT, or broadcast frame i.e. be the same as the current value of the receive frame sequence counter. If this is not the case it indicates that a frame that has been sent has not been received, so a NACK is sent to indicate the last correctly received data frame, FCT or broadcast frame.

NOTE When an idle frame is received that has the correct frame sequence number there is no need to send an ACK. The correct frame sequence number indicates that the last data frame, FCT or broadcast frame has been received correctly. If the ACK sent for this frame has gone missing, it does not matter unless the frame retry buffer fills up which will not happen unless more data frames, FCTs or broadcast frames are sent. In which case other ACKs will be sent.

b. When requested or pending NACKs shall be sent in the error state of the Retry layer error state machine, as soon as possible, taking into account control word precedence (see section 5.3.4).

c. When an NACK is sent it shall contain the current value of the receive frame sequence counter with the receive polarity flag inverted.

NOTE If an error occurs when the Retry layer error state machine is in valid state the receive polarity flag of the receive frame sequence counter is inverted before a NACK is requested to be sent. Therefore, in this particular case, the polarity flag of this NACK sent will match the polarity flag of the last valid data frame received.

d. When an ACK is requested to be sent, any pending NACK shall be cancelled and not sent.

NOTE A pending NACK is one that has been requested to be sent but was not able to be sent immediately because control or data words with higher precedence were being sent.

5.8.9.3 Retry layer error state machine

a. The retry layer state machine shall be used to determine when the polarity flag of the receive frame sequence counter shall be inverted provided that the clause 5.8.6.2r.2 is not fulfilled.

NOTE The state diagram for the retry layer error state machine is illustrated in Figure 5-7

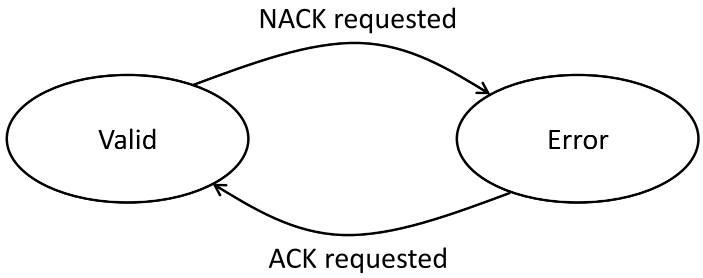

Figure 5-7 Retry layer error state machine

5.8.9.3.1 Valid State

a. The Valid state shall be entered on the following conditions:

   1. ColdReset command.
   2. WarmReset command.
   3. Remote Flush.
   4. From the Error state, when an ACK request is received.

b. When in the Valid state the retry layer state machine shall initiate the following action:

   1. When a NACK request is received, the polarity flag of the receive frame sequence counter shall be inverted.

c. The Retry layer error state machine shall leave the Valid state on the following conditions:

   1. When a NACK request is received, move to the Error state.

   NOTE The Valid state is summarised in Table 5-32.

Table 5-17 Valid State

State	Valid
Entry	ColdReset command. WarmReset command. Remote Flush. From the Error state, when an ACK request is received
Action	When a NACK request is received, the polarity flag of the receive frame sequence counter shall be inverted
Exit	When a NACK request is received, move to the Error state.

5.8.9.3.2 Error State

a. The Error state shall be entered on the following condition:

   1. From the Valid state, when a NACK request is received.

b. The retry layer error state machine shall leave the Error state on the following conditions:

1. When an ACK request is received, move to the Valid state.

NOTE The Error state is summarised in Table 5-33.

Table 5-18 Error State	
State	Error
Entry	From the Valid state, when a NACK request is received
Action	No action
Exit	When an ACK request is received, move to the Valid state.

5.8.9.4 Receiving ACKs

a. When an ACK is received it shall be checked for errors using its CRC.

b. An ACK is considered valid when it has a valid CRC and the polarity flag of its frame sequence number has the same value as the polarity flag of the transmit frame sequence counter.

c. When a valid ACK is received, all data frames, broadcast frames, and FCTs sent with a frame sequence number less than or equal to the frame sequence number in the ACK shall be deleted from the frame retry buffer, taking into account counter roll over effects (i.e. calculation is modulo 128).

5.8.9.5 Receiving NACKs

a. When a NACK is received it shall be checked for errors using its CRC.

b. A NACK is considered valid when it has a valid CRC and the polarity flag of its frame sequence number has the same value as the polarity flag of the transmit frame sequence counter.

c. When a valid NACK is received the following actions shall be executed in the specified order:

1. All data frames, FCTs and broadcast frames, sent with a frame sequence number less than or equal to the frame sequence number in the NACK shall be deleted from the frame retry buffers, taking into account counter roll over effects (i.e. calculation is modulo 128).

2. The transmit frame sequence counter is set to the value of the NACK sequence number.

3. The polarity flag of the transmit frame sequence counter shall be inverted.

4. Any record of the frame sequences sent related to the content of the retry buffers shall be cleared.

NOTE When the frames/FCTs are resent they will be given new sequence numbers.

5. A RETRY word shall be sent before the contents of the retry buffer are sent

6. The contents of the retry buffers shall be sent following the precedence indicated in section 5.3.4.

5.9 Multi-lane layer

5.9.1 Multi-lane layer service interface

The service primitives that shall be associated with the multi-lane layer service are:

TX_WORD.request;

TX_WORD.indication;

RX_WORD.indication;

5.9.1.1 TX_WORD.request

5.9.1.1.1 Function

The retry layer shall pass a TX_WORD.request primitive to the multi-lane layer to send a data word or control word that forms part of a data frame, broadcast frame, idle frame, FCT, ACK or NACK over the SpaceFibre link.

5.9.1.1.2 Semantics

The TX_WORD.request primitive shall provide the following parameters:

TX_WORD.request (Word)

5.9.1.1.3 When Generated

The TX_WORD.request primitive shall be passed to the multi-lane layer when the retry layer has information to send over the SpaceFibre link.

5.9.1.1.4 Effect On Receipt

On receipt of the TX_WORD.request primitive the multi-lane layer shall send the data word or control word over the SpaceFibre interface.

5.9.1.2 TX_WORD. indication

5.9.1.2.1 Function

The multi-lane layer shall pass a TX_WORD.indication primitive to the retry layer to indicate that it is ready to accept a new data or control word.

5.9.1.2.2 Semantics

The TX_WORD.indication primitive shall have no parameters.

5.9.1.2.3 When Generated

The TX_WORD.indication primitive shall be passed to the Retry Layer when the Multi-lane layer can accept a new request to transfer a data or control word.

5.9.1.2.4 Effect On Receipt

On receipt of the TX_WORD.indication primitive the retry layer shall send a new TX_WORD.request as soon as there is another data or control word ready to send.

5.9.1.3 RX_WORD.indication

5.9.1.3.1 Function

The multi-lane layer shall pass a RX_WORD.indication primitive to the retry layer to indicate that a data word or control word has been received.

5.9.1.3.2 Semantics

The RX_WORD.indication primitive provides parameters as follows:

RX_WORD.indication (Word).

5.9.1.3.3 When Generated

The RX_WORD.indication primitive shall be passed to the retry layer when a data word or control word is received.

5.9.1.3.4 Effect On Receipt

The effect on receipt of the RX_WORD.indication primitive by the retry layer shall result in the word being accepted by the retry layer.

5.9.2 Lane control

a. The Required Number of Lanes management parameter (required lanes) shall specify the number of lanes that are to be used to form the SpaceFibre link.

b. The number of lanes that are active and ready to send data and control words (active lanes) shall be indicated from the lane layer to the multi-lane layer.

c. If the number of required lanes is less than the number of active lanes, the number of used lanes shall be set to the number of required lanes.

d. If the number of required lanes is the same as the number of active lanes, the number of used lanes shall be set to the number of required lanes.

e. If the number of required lanes is more than the number of active lanes, the number of used lanes shall be set to the number of active lanes.

f. Each lane shall be given a lane number, starting at 1 and incrementing for each physical link.

NOTE For example, if there are four possible lanes they would be given lane numbers 1, 2, 3 and 4.

5.9.3 Lane synchronisation

a. When the number of required lanes or the number of active lanes changes, lane synchronisation shall occur.

b. During lane synchronisation the number of used lanes shall be determined.

c. Lanes shall be assigned lane umbers as follows: the lowest number lane that is active is assigned lane number 1, the next lowest number lane that is active is assigned lane number 2, and subsequent lanes are assigned lane numbers in this way until enough lanes have been assigned lane numbers to cover the number of used lanes.

d. All other physical links that are active shall be assigned a null lane number (zero).

e. An LSYNC control word shall be sent over each of the active lanes containing the lane number assigned to that lane.

f. The LSYNC control word shall be sent over each of the active lanes at approximately the same time, i.e. within the time it takes to transmit one control word (TBC).

g. Once the LSYNC control word has been sent, data and control words can be distributed over the used lanes for transmission.

h. When an LSYNC control word is received, the lane synchroniser shall wait for LSYNC control words to be received over each of the other active lanes.

i. When LSYNC control words have been received on all active lanes, the lane concentrator shall be updated with the lanes it is to concentrate data from.

j. The active lanes that receive an LSYNC control word containing a null (zero) shall not be included in the lane concentration.

k. The active lanes that receive an LSYNC control word containing a lane number shall be included in the lane concentration.

l. If an LSYNC control word is not received on a lane when LSYNC control words are received on the other active lanes, all the lanes are instructed to reinitialise.

NOTE Link re-initialisation will automatically invoke resynchronisation of the lanes.

m. If the received LSYNC control words do not form a proper integer series (i.e. 1, 2, 3, etc) of lane numbers with no duplicate lane numbers, other than possible duplicate null lanes, and no missing lane numbers, all the lanes are instructed to reinitialise.

NOTE This is likely to be because it was lost and replaced by an RXERR control word or the result of a serious error in the transmitter, which re-initialisation might not correct.

n. If after invoking re-initialisation of lanes, lane synchronisation still fails, the error shall be flagged as a permanent error and the user application informed.

5.9.4 Lane distribution

a. Data and control words to be sent over the SpaceFibre link shall be distributed over the used lanes.

b. The first data or control word shall be sent over the lowest number used lane, the next data or control word over the next lowest number used lane, and so on with data or control words being sent over each of the used lanes.

c. Data and control words shall not be sent over active lanes which are not used lanes.

NOTE Such lanes are providing hot standby and will be sending IDLE control words.

5.9.5 Lane concentration

a. When a data or control word arrives over a lane, it shall be placed in a small synchronisation FIFO.

b. There shall be one synchronisation FIFO for each lane.

c. The synchronisation FIFO shall be able to store at least three data or control words.

d. Data or control words shall be read out of the synchronisation FIFOs for each used lane in lane number order.

e. Following lane synchronisation, the first data or control word shall be read out of the lane with the lowest lane number, the next data or control word over the next lowest number used lane, and so on with data and control words being read out of the synchronisation FIFOs of each of the used lanes.

NOTE The lane numbers of the concentrator can be different to that of the distributor.

5.9.6 Lane selection

It shall be possible to switch off or bypass multi-lane operation so that the SpaceFibre link operates over one lane only.

5.10 Lane layer

5.10.1 Lane layer service interface

The service primitives that shall be associated with the lane layer service are:

TX_WORD.request;

TX_WORD.indication;

RX_WORD.indication;

LINK_STATE.indication.

5.10.1.1 TX_WORD.request

5.10.1.1.1 Function

The multi-lane layer shall pass a TX_WORD.request primitive to the lane layer to send a data word or control word that forms part of a data frame, broadcast frame, idle frame, FCT, ACK or NACK over the SpaceFibre link.

5.10.1.1.2 Semantics

The TX_WORD.request primitive shall provide the following parameters:

TX_WORD.request (Word)

5.10.1.1.3 When Generated

The TX_WORD.request primitive shall be passed to the lane layer when the multi-lane layer has information to send over the SpaceFibre link.

5.10.1.1.4 Effect On Receipt

On receipt of the TX_WORD.request primitive the lane layer shall send the data word or control word over the SpaceFibre link.

5.10.1.2 TX_WORD.indication

5.10.1.2.1 Function

The lane layer shall pass a TX_WORD.indication primitive to the multi-lane layer to indicate that it is ready to accept a new data or control word.

5.10.1.2.2 Semantics

The TX_WORD.indication primitive shall have no parameters.

5.10.1.2.3 When Generated

The TX_WORD.indication primitive shall be passed to the multi-lane layer when the lane layer can accept a new request to transfer a data or control word.

5.10.1.2.4 Effect On Receipt

On receipt of the TX_WORD.indication primitive the multi-lane layer shall send a new TX_WORD.request as soon as there is another data or control word ready to send.

5.10.1.3 RX_WORD.indication

5.10.1.3.1 Function

The lane layer shall pass a RX_WORD.indication primitive to the multi-lane layer to indicate that a data word or control word has been received.

5.10.1.3.2 Semantics

The RX_WORD.indication primitive provides parameters as follows:

RX_ WORD.indication (Word).

5.10.1.3.3 When Generated

The RX_ WORD.indication primitive shall be passed to the multi-lane layer when a data word or control word is received.

5.10.1.3.4 Effect On Receipt

The effect on receipt of the RX_ WORD.indication primitive by the multi-lane layer shall be that the received word is taken by the multi-lane layer and interleaved with words from other lanes.

5.10.1.4 LANE_STATUS.indication

5.10.1.4.1 Function

The lane layer shall pass a LANE_STATUS.indication primitive to the multi-lane layer to indicate that the status of the lane has changed.

5.10.1.4.2 Semantics

The LANE_STATUS.indication primitive provides parameters as follows:

LANE_STATUS.indication (Lane Status).

NOTE The lane status values are TBD.

5.10.1.4.3 When Generated

The LANE_STATUS.indication primitive shall be passed to the multi-lane layer when the lane status has changed.

5.10.1.4.4 Effect On Receipt

On receipt of the LANE_STATUS.indication primitive the multi-lane layer shall react depending on the lane status. If the lane status is lane not ready, the multi-lane layer will stop sending words to the lane layer and cease reading words from that layer. The multi-lane layer will also resynchronise the remaining active lanes. If the lane status is lane ready, the multi-lane layer will resynchronise the active lanes and start sending words to the lane layer and reading words from that layer.

5.10.2 Lane initialisation and standby management

a. The Lane Initialisation state machine shall control the initialisation of the SpaceFibre link: establishing the connection and responding to standby management requests.

NOTE The state diagram for the lane initialisation state machine is illustrated in Figure 5-8.

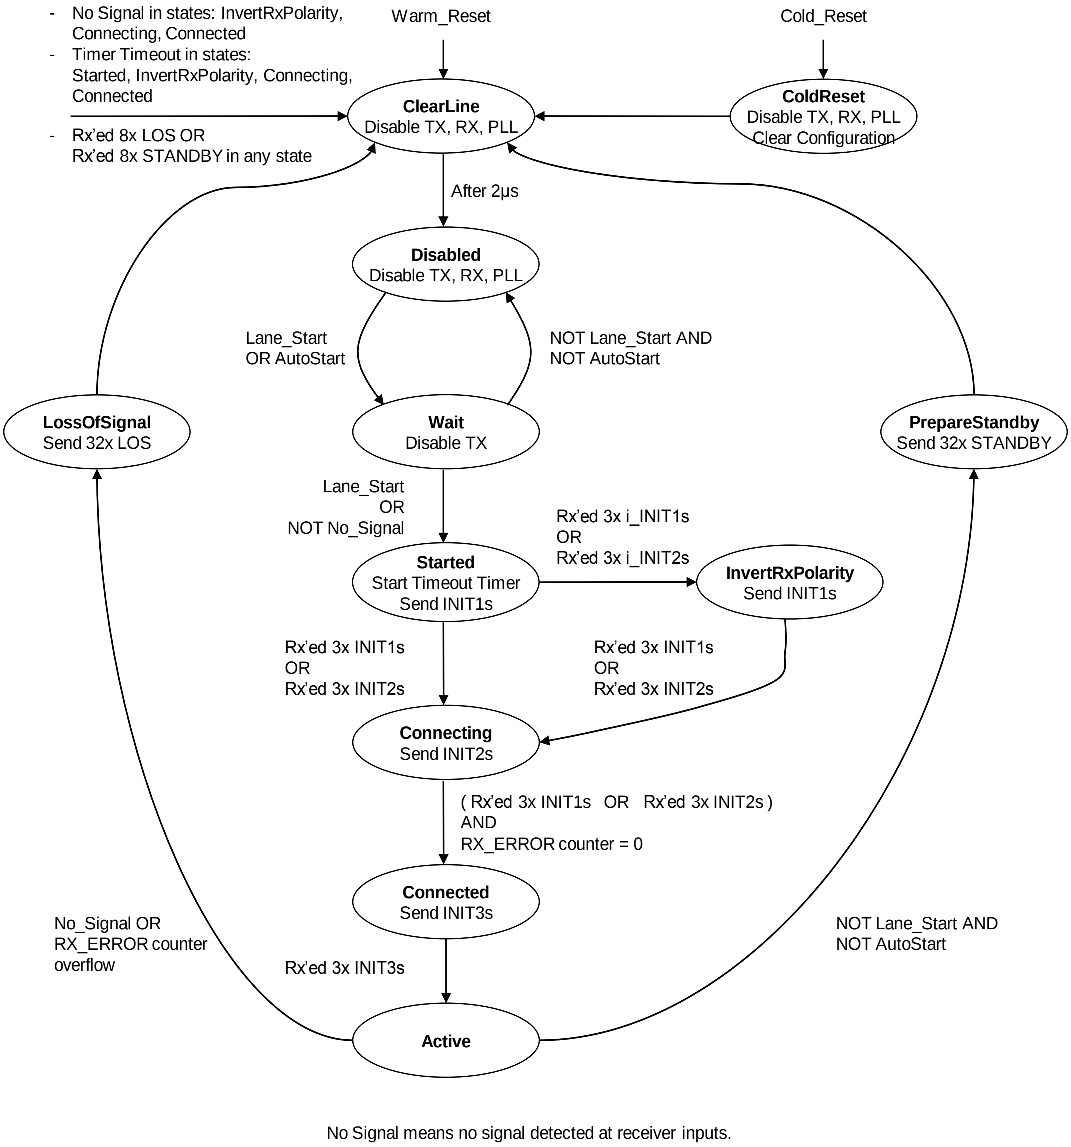

Figure 5-8 Lane Initialisation State Machine

5.10.2.1 RX_ERROR counter

NOTE Various alternatives to the clock-data recovery lock detection are being considered. The RX_ERROR counter operation will be finalised once this research has been concluded.

a. An RX_ERROR counter shall be used to check whether the clock-data recover circuitry has locked and is performing bit and symbol synchronisation correctly.

b. The initial value of the RX_ERROR counter shall be 8 TBC.

c. Every time a receive error is detected the RX_ERROR counter shall be incremented.

d. Every 32 (TBC) received words the RX_ERROR counter shall be decremented.

e. If the RX_ERROR counter reaches zero, the clock-data recovery circuitry shall be assumed to be correctly locked.

NOTE This can only result when 32x8 = (TBC)

5.10.2.2 ColdReset state

a. The Lane Initialisation state machine shall enter the ColdReset state on one of the following conditions:

1. Power on reset.
2. ColdReset command.

NOTE Commands are issued via the Network Management interface which sets appropriate bits in control registers.

b. When in the ColdReset state the Lane Initialisation state machine shall initiate the following actions:

1. Reset of the SpaceFibre Lane.
2. Reset of all the Lane Configuration registers.
3. Assert the Cold_Reset_Flag.

NOTE This flag is used to indicate that it was a cold reset that caused the re-initialisation.

4. Disable the transmitter driver, receiver and related transmit PLL and receive clock recovery.

NOTE This is to save power.

c. The Lane Initialisation state machine shall leave the ColdReset state on one of the following conditions:

1. Unconditionally, move to the ClearLine state.

d. The ColdReset state is summarised in Table 5-19.

Table 5-19 ColdReset State

State	ColdReset
Entry	Power on reset. ColdReset command.
Action	Reset of the SpaceFibre Lane. Reset of all the Lane Configuration registers. Disable the transmitter driver, receiver and related transmit PLL and receive clock recovery.
Exit	Unconditionally, move to the ClearLine state.

5.10.2.3 ClearLine state

a. The Lane Initialisation state machine shall enter the ClearLine state on one of the following conditions:

1. From the ColdReset state, unconditionally.
2. On WarmReset command.
3. From the Started state, when the initialisation timeout timer expires.

    NOTE The initialisation timeout timer is started in the Start state and stopped in the Active state. It will timeout if initialisation takes longer than 20 μs. This period allows plenty of time for the receive clock recovery at the far end of the lane to lock onto the transmitted signal, a response to be generated, returned to the near end and the local receiver to lock onto the returned signal.

4. From the InvertRxPolarity state, when the initialisation timeout timer expires.
5. From the Connecting state, when the initialisation timeout timer expires.
6. From the Connected state, when the initialisation timeout timer expires.
7. From any state when 8 consecutive LOS words are received.
8. From any state when 8 consecutive STANDBY words are received.
9. From the InvertRxPolarity state, when No Signal is detected at the receiver.
10. From the Connected state, when No Signal is detected at the receiver.
11. From the Connecting state, when No Signal is detected at the receiver.
12. From the PrepareStandby state, after sending 32 STANDBY control words.
13. From the LossOfSignal state, after sending 32 LOST_SIGNAL control words.

b. When in the ClearLine state the Lane Initialisation state machine shall initiate the following actions:

1. Start a 2 μs timeout timer.
2. Disable the transmitter driver, receiver and related transmit PLL and receive clock recovery.
3. Do NOT reset the Lane Configuration registers.
4. Switch off receiver bit inversion.

c. The Lane Initialisation state machine shall leave the ClearLine state on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.
2. When the 2 μs timeout timer expires, move to the Disabled state.

   NOTE This is more than long enough for signals to propagate through the longest permitted fibre optic cable (100m) and back again, and to allow time for the signals to be processed at each end of the cable. The result is that the lanes forming the link will be completely reset at the near end of the link and at least warm reset at the far end.

d. The ClearLine state is summarised in Table 5-20.

Table 5-20 ClearLine State

State	ClearLine
Entry	From the ColdReset state, unconditionally. On WarmReset command. From the Started state, when the initialisation timeout timer expires. From the InvertRxPolarity state, when the initialisation timeout timer expires. From the Connecting state, when the initialisation timeout timer expires. From the Connected state, when the initialisation timeout timer expires. From any state when 8 LOS words are received. From any state when 8 consecutive STANDBY words are received. From the InvertRxPolarity state, when No Signal is detected at the receiver. From the Connected state, when No Signal is detected at the receiver. From the Connecting state, when No Signal is detected at the receiver From the PrepareStandby state, after sending 32 STANDBY control words. From the LossOfSignal state, after sending 32 LOST_SIGNAL control words.
Action	Start a 2 μs timeout timer. Disable the transmitter driver, receiver and related transmit PLL and receive clock recovery. Do NOT reset the Lane Configuration registers. Switch off receiver bit inversion.
Exit	When ColdReset is asserted, move to the ColdReset state. When the 2 μs timeout timer expires, move to the Disabled

	state.

5.10.2.4 Disabled state

a. The Lane Initialisation state machine shall enter the Disabled state on the following condition:

   1. From the ClearLine state, after waiting there for 2μs.

b. When in the Disabled state the Lane Initialisation state machine shall initiate the following action:

   1. Disable the transmitter driver, receiver and related transmit PLL and receive clock recovery.

c. The Lane Initialisation state machine shall leave the Disabled state on one of the following conditions, which are to be evaluated in the order given:

   1. When ColdReset is asserted, move to the ColdReset state.
   2. When WarmReset is asserted, move to the ClearLine state.
   3. When Lane_Start OR AutoStart is asserted, move to the Wait state.
   4. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

d. The Disabled state is summarised in Table 5-21.

Table 5-21 Disabled State

State	Disabled
Entry	From the ClearLine state, after waiting there for 2μs.
Action	Disable the transmitter driver, receiver and related transmit PLL and receive clock recovery.
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When Lane_Start OR AutoStart is asserted, move to the Wait state. When 8 consecutive LOST_SIGNAL OR 8 consecutive STANDBY control words have been received, move to the ClearLine state.

5.10.2.5 Wait state

a. The Lane Initialisation state machine shall enter the Wait State on one of the following conditions:

   1. From the Disabled state, when Lane_Start OR AutoStart is asserted.

b. When in the Wait State the Lane Initialisation state machine shall initiate the following actions:

1. Disable the transmitter driver and related transmit PLLs.
2. Enable the receiver, and receiver clock recovery.
3. Enable reception of Lane Layer Control Words (LLCW).

c. The Lane Initialisation state machine shall leave the Wait State on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.
2. When WarmReset is asserted, move to the ClearLine state.
3. When Lane_Start AND AutoStart are both deasserted, move to the Disabled state.
4. When Lane_Start is asserted OR No_Signal is deasserted, move to the Started state.
5. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

d. The Wait State is summarised in Table 5-26.

Table 5-22 Wait State

State	Wait
Entry	From the Disabled state, when Lane_Start OR AutoStart is asserted.
Action	Disable the transmitter driver and related transmit PLLs. Enable the receiver, and receiver clock recovery. Enable reception of Lane Layer Control Words (LLCW).
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When Lane_Start AND AutoStart are both deasserted, move to the Disabled state. When Lane_Start is asserted OR No_Signal is deasserted, move to the Started state. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

5.10.2.6 Started state

a. The Lane Initialisation state machine shall enter the Start state on the following condition:

1. From the Wait state, when Lane_Start is asserted or No_Signal is deasserted.

b. When in the Start state the Lane Initialisation state machine shall initiate the following actions:

1. Start a 20 μs initialisation timeout timer, on entry to the state.

NOTE 20 μs allows plenty of time for the receive clock recovery at the far end of the lane to lock onto the transmitted signal, a response to be generated and returned to the near end.

2. Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery.

3. Send INIT1 control words.

4. Receive lane layer control words (LLCW).

5. Set the RXERR words counter to TBC value.

NOTE The Initial Value for RXERR Words Counter is a generic or a management parameter, whose value determines the minimum error rate acceptable for the PLL to be considered locked.

c. The Lane Initialisation state machine shall leave the Start state on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.

2. When WarmReset is asserted, move to the ClearLine state.

3. When three INIT1 or three INIT2 control words are received without any intervening RXERR control words, move to the Connecting state.

4. When three inverse INIT1 or three inverse INIT2 control words are received without any intervening RXERR control words, move to the InvertRxPolarity state.

5. When the initialisation timeout timer expires, move to the ClearLine state.

6. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

d. The Start state is summarised in Table 5-23.

Table 5-23 Started State

State	Started
Entry	From the Wait state, when Lane_Start is asserted and No_signal is de-asserted.
Action	Start a 20 μs initialisation timeout timer, on entry to the state. Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery. Send INIT1 control words. Receive lane layer control words (LLCW). Set the RXERR words counter to TBC value.
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When three INIT1 or three INIT2 control words are received without any intervening RXERR control words, move to the Connecting state.

	When three inverse INIT1 or three inverse INIT2 control words are received without any intervening RXERR control words, move to the InvertRxPolarity state. When the initialisation timeout timer expires, move to the ClearLine state. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

5.10.2.7 InvertRxPolarity state

a. The Lane Initialisation state machine shall enter the InvertRxPolarity state on one of the following conditions:

1. From the Started state, when three inverse INIT1 control words are received without any intervening RXERR control words.
2. From the Started state, when three inverse INIT2 control words are received without any intervening RXERR control words.

b. When in the InvertRxPolarity state the Lane Initialisation state machine shall initiate the following actions:

1. Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery.
2. Switch on receiver bit inversion.
3. Receive lane layer control words (LLCW).
4. Set the RXERR words counter to TBC value.

c. The Lane Initialisation state machine shall leave the InvertRxPolarity state on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.
2. When WarmReset is asserted, move to the ClearLine state.
3. When No Signal is detected at the receiver inputs, move to the ClearLine state.
4. When three INIT1 control words are received without any intervening RXERR control words, move to the Connecting state.
5. When three INIT2 control words are received without any intervening RXERR control words, move to the Connecting state.
6. When the initialisation timeout timer expires, move to the ClearLine state.
7. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

d. The StartPolarity state is summarised in Table 5-24.

Table 5-24 InvertRxPolarity State

State	InvertRxPolarity

Entry	From the Started state, when three inverse INIT1 control words are received without any intervening RXERR control words. From the Started state, when three inverse INIT2 control words are received without any intervening RXERR control words.
Action	Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery. Switch on receiver bit inversion. Receive lane layer control words (LLCW). Set the RXERR words counter to TBC value.
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When No Signal is detected at the receiver inputs, move to the ClearLine state. When three INIT1 control words are received without any intervening RXERR control words, move to the Connecting state. When three INIT2 control words are received without any intervening RXERR control words, move to the Connecting state. When the initialisation timeout timer expires, move to the ClearLine state. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

5.10.2.8 Connecting state

a. The Lane Initialisation state machine shall enter the Connecting state on one of the following conditions:

   1. From the Started state, when three INIT1 control words are received without any intervening RXERR control words.

   2. From the Started state, when three INIT2 control words are received without any intervening RXERR control words.

   3. From the InvertRxPolarity state, when three INIT1 control words are received without any intervening RXERR control words.

   4. From the InvertRxPolarity state, when three INIT2 control words are received without any intervening RXERR control words.

b. When in the Connecting state the Lane Initialisation state machine shall initiate the following actions:

   1. Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery.

   2. Send INIT2 control words.

   3. Receive lane layer control words (LLCW).

4. Increase the RXERR words counter when a RXERR is received.

5. Decrease the RXERR words counter by one every 32 words.

c. The Lane Initialisation state machine shall leave the Connecting state on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.

2. When WarmReset is asserted, move to the ClearLine state.

3. When No Signal is detected at the receiver inputs, move to the ClearLine state.

4. When three INIT2 control words are received with the RXERR words counter being zero, move to the Connected state.

5. When three INIT3 control words are received with the same initialisation parameters and with the RXERR words counter being zero, move to the Connected state

6. When the initialisation timeout timer expires, move to the ClearLine state.

7. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

d. The Connecting state is summarised in Table 5-26.

Table 5-25 Connecting State

State	Connecting
Entry	From the Started state, when three INIT1 control words are received without any intervening RXERR control words. From the Started state, when three INIT2 control words are received without any intervening RXERR control words. From the InvertRxPolarity state, when three INIT1 control words are received without any intervening RXERR control words. From the InvertRxPolarity state, when three INIT2 control words are received without any intervening RXERR control words.
Action	Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery. Send INIT2 control words. Receive lane layer control words (LLCW) Increase the RXERR words counter when a RXERR is received. Decrease the RXERR words counter by one every TBC words.
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When No Signal is detected at the receiver inputs, move to the ClearLine state. When three INIT2 control words are received with the RXERR

	words counter being zero, move to the Connected state. When three INIT3 control words are received with the same initialisation parameters and with the RXERR words counter being zero, move to the Connected state. When the initialisation timeout timer expires, move to the ClearLine state. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

5.10.2.9 Connected state

a. The Lane Initialisation state machine shall enter the Connected state on one of the following conditions:

1. From the Connecting state, when three INIT2 control words are received.
2. From the Connecting state, when three INIT3 control words are received with the same initialisation parameters.

b. When in the Connected state the Lane Initialisation state machine shall initiate the following actions:

1. Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery.
2. Send INIT3 control words with the initialisation parameters set in the most significant byte as follows:
   (a) Bit 0: the Remote_Flush flag,
   (b) Bit 1: the Lane_Start flag,
   (c) Bit 2: the Data_Scrambled flag,
   (d) Bits 3 to 7: reserved bits that are set to zero when transmitted and ignored when received.
3. Receive lane layer control words (LLCW).
4. Store the initialisation parameters supplied by the INIT3 control word most significant value.
5. When leaving this state deassert the Cold_Reset_Flag.

   NOTE This flag is used to indicate that it was a cold reset that caused the re-initialisation.
6. When leaving this state apply the initialisation changes in the initialisation register:
   (a) Assert the Remote Flush condition if the Remote_Flush_Flag is set and the Cold_Reset_Flag is not set.
   (b) Pass the Lane_Start value to the management layer.
   (c) Pass the Data_Scrambled value to the framing layer.
7. Increase the RXERR words counter when a RXERR is received.

8. Decrease the RXERR words counter by one every TBC words.

c. The Lane Initialisation state machine shall leave the Connected state on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.
2. When WarmReset is asserted, move to the ClearLine state.
3. When No Signal is detected at the receiver inputs, move to the ClearLine state.
4. When three INIT3 control words are received with the same initialisation parameters and with the RXERR words counter being zero, move to the Active state.
5. When the initialisation timeout timer expires, move to the ClearLine state.
6. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state

d. The Connected state is summarised in Table 5-26.

Table 5-26 Connected State

State	Connected
Entry	From the Connecting state, when three INIT2 control words are received without intervening RXERR control words. From the Connecting state, when three INIT3 control words are received with the same initialisation parameters without intervening RXERR control words.
Action	Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery. Send INIT3 control words. Receive lane layer control words (LLCW). Store the initialisation parameters supplied by the INIT3 control word in a register, and apply the new configuration if required. When leaving this state apply the initialisation changes in the initialisation register. Increase the RXERR words counter when a RXERR is received. Decrease the RXERR words counter by one every TBC words.
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When No Signal is detected at the receiver inputs, move to the ClearLine state. When three INIT3 control words are received with the same initialisation parameters and with the RXERR words counter being zero, move to the Active state. When the initialisation timeout timer expires, move to the ClearLine state.

	When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state

5.10.2.10 Active state

a. The Lane Initialisation state machine shall enter the Active state on the following conditions:

1. From the Connected state, when three INIT3 control words with the same capability field value have been received without intervening RXERR control words.

   NOTE The Lane Initialisation state machine enters the Active state when initialisation is complete.

b. When in the Active state the Lane Initialisation state machine shall initiate the following actions:

1. Stop the initialisation timeout timer.
2. Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery.
3. Enable transmission of data and control words from the lane control or retry layer.
4. Pass received data and control words up to the lane control or retry layer.
5. Increase the RXERR words counter when a RXERR is received.
6. Decrease the RXERR words counter by one every RXERR counter decrement period words, which is a configuration parameter.
7. Filter out Lane control words (see Table 5-1) so that they are not passed up to the lane control or retry layer.

   NOTE They are used by the Lane Initialisation state machine and nothing is passed up to the multi-lane layer or retry layer in their place.

c. The Lane Initialisation state machine shall leave the Active state on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.
2. When WarmReset is asserted, move to the ClearLine state.
3. When No Signal at Receiver signal is asserted indicating that there is a loss of signal at the receiver, move to the LossOfSignal state.
4. When the RX_ERR words counter is TBC value, move to the LossOfSignal state.

   NOTE If the RXERR words error counter overflows there is a persistent error and the lane should be reinitialized. The multi-lane layer may disable a lane when a persistent error occurs and use instead another redundant lane.

5. When the Lane_Start and Auto_Start signals are both de-asserted, move to the PrepareStandby state.

6. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

d. The Active state is summarised in Table 5-27 Active State.

Table 5-27 Active State

State	Active
Entry	From the Connected state, when three INIT3 control words with the same initialisation parameters have been received without intervening RXERR control words..
Action	Stop the initialisation timeout timer. Enable the transmitter, transmitter PLL, receiver, and receiver clock recovery. Enable transmission of data and control words from the lane control or retry layer. Pass received data and control words up to the lane control or retry layer. Filter out Lane control words (see Table 5-1) so that they are not passed up to the lane control or retry layer.
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When No Signal at Receiver signal is asserted indicating that there is a loss of signal at the receiver, move to the LossOfSignal state. When the RX_ERR words counter is TBC value, move to the LossOfSignal state. When the Lane_Start and Auto_Start signals are both de-asserted, move to the PrepareStandby state. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

5.10.2.11 PrepareStandby state

a. The Lane Initialisation state machine shall enter the PrepareStandby state on the following condition:

1. From the Active state, when Lane_Start and AutoStart signals are both de-asserted.

b. When in the PrepareStandby state the Lane Initialisation state machine shall initiate the following action:

1. Send 32 STANDBY control words.

NOTE 32 STANBY control words are sufficient to make sure that the first STANDBY control word has passed through the receive pipeline and receive elastic buffer at the far end of the lane, before the subsequent No Signal at receiver, caused by the near end of the lane entering the ClearLine state, is detected by the far end receiver.

c. The Lane Initialisation state machine shall leave the PrepareStandby state on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.
2. When WarmReset is asserted, move to the ClearLine state.
3. When 32 STANDBY control words have been sent, move to the ClearLine state.
4. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

d. The PrepareStandby state is summarised in Table 5-28.

Table 5-28 PrepareStandby State

State	PrepareStandby
Entry	From the Active state, when Lane_Start and Auto_Start signals are both de-asserted.
Action	Send 32 STANDBY control words.
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When 32 STANDBY control words have been sent, move to the ClearLine state. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

5.10.2.12 LossOfSignal state

a. The Lane Initialisation state machine shall enter the LossOfSignal state on the following conditions:

1. From the Active state when No Signal at Receiver signal is asserted, indicating that there is no signal present on the receiver inputs.
2. From the Active state when the RX_ERR words counter overflows.

NOTE A counter increases its value every time a RX_ERR word is received, and periodically decreases its value (leaky-bucket counter). The overflow of this counter causes the output condition from Active state and an indication of a persistent error.

b. When in the LossOfSignal state the Lane Initialisation state machine shall initiate the following actions:

1. Send 32 LOST_SIGNAL control words.
2. Pass at least one (RXERR) word up to the multi-lane or retry layer.

c. The Lane Initialisation state machine shall leave the LossOfSignal state on one of the following conditions, which are to be evaluated in the order given:

1. When ColdReset is asserted, move to the ColdReset state.
2. When WarmReset is asserted, move to the ClearLine state.
3. When 32 LOST_SIGNAL control words have been sent, move to the ClearLine state.
4. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

d. The LossOfSignal state is summarised in Table 5-29.

Table 5-29 LossOfSignal State

State	LossOfSignal
Entry	From the Active state when No Signal at Receiver signal is asserted, indicating that there is no signal present on the receiver inputs. From the Active state when the RX_ERR words counter overflows.
Action	Send 32 LOST_SIGNAL control words. Pass at least one (RXERR) word up to the multi-lane or retry layer.
Exit	When ColdReset is asserted, move to the ColdReset state. When WarmReset is asserted, move to the ClearLine state. When 32 LOST_SIGNAL control words have been sent, move to the ClearLine state. When 8 consecutive LOST_SIGNAL or 8 consecutive STANDBY control words have been received, move to the ClearLine state.

5.10.3 Data rate adjustment

a. The transmitters at each end of a lane shall send data at the same nominal data signalling rate.

b. The maximum permitted difference in the data signalling rates, when each end of the lane is operating at nominally the same data signalling rate (e.g. 2 Gbits/s) shall be 100 parts per million.

NOTE The data signalling rate of a transmitter at one end of a lane and a receiver at the other end of a lane might be different due to differences in the local clocks being used at each end of the lane.

c. A receive elastic buffer shall be used to compensate for differences in the data signalling rate at each end of the lane.

d. Skip control words shall be inserted regularly in the transmit data stream.

NOTE When a SKIP arrives at the receiver at the other end of the lane it might have been inverted by PCB layout to form an inverse SKIP control word.

e. A skip control word shall be sent every 5,000 words.

f. When a skip or inverse skip control word is read from a receive elastic buffer that is more than half full, the skip or inverse skip control word shall be removed from the receive elastic buffer and the following character shall be read from the buffer.

g. When a skip or inverse skip control word is read from a receive elastic buffer that is less than half full, the skip or inverse skip control word shall be left in the buffer, so that it can be read again the next time the buffer is read.

h. A particular skip or inverse skip control word shall be read at most twice, i.e. the second time it is read it is removed from the receive elastic buffer.

5.10.4 Idle words

a. When there is no other data or control word to send an idle word shall be sent.

NOTE It is essential not to have gaps in the data being sent because the PLL in the receiver has to maintain lock on the incoming data stream.

b. When received, idle or inverse idle control words shall be discarded when they are read out of the receive elastic buffer.

NOTE Although only IDLEs are sent they might be inverted into inverse IDLEs by the time they reach the receiver.

5.10.5 Parallel loopback

a. A parallel loopback facility should be provided in the SpaceFibre CODEC for test purposes.

NOTE This facility is optional.

b. When enabled, the parallel loopback shall pass the stream of words to be transmitted from the lane layer to the receive input of the lane layer.

NOTE This connects the output of the transmit side of the lane layer to the input of the receive side of the lane layer as illustrated in Figure 4-2.

5.11 Encoding layer

5.11.1 Encoding layer service interface

The service primitives that shall be associated with the encoding layer service are:

TX_WORD.request;

RX_WORD.indication;

SIGNAL.indication;

RX_INVERSION.request.

5.11.1.1 TX_WORD.request

5.11.1.1.1 Function

The lane layer shall pass a TX_WORD.request primitive to the encoding layer to send a data word or control word over the lane.

5.11.1.1.2 Semantics

The TX_WORD.request primitive shall provide the following parameters:

TX_WORD.request (Word)

5.11.1.1.3 When Generated

The TX_WORD.request primitive shall be passed to the encoding layer when the lane layer has information to send over the lane.

5.11.1.1.4 Effect On Receipt

On receipt of the TX_WORD.request primitive, the encoding layer shall encode the data word or control and send it over the lane.

5.11.1.2 RX_WORD.indication

5.11.1.2.1 Function

The encoding layer shall pass a RX_WORD.indication primitive to the lane layer to indicate that a data word has been received.

5.11.1.2.2 Semantics

The RX_WORD.indication primitive provides parameters as follows:

RX_ WORD.indication (Word).

5.11.1.2.3 When Generated

The RX_ WORD.indication primitive shall be passed to the lane layer when a data or control word is received.

5.11.1.2.4 Effect On Receipt

The effect on receipt of the RX_ WORD.indication primitive by the lane layer shall be that the received word is placed in the receive elastic buffer.

5.11.1.3 SIGNAL.indication

5.11.1.3.1 Function

The encoding layer shall pass a SIGNAL.indication primitive to the lane layer to indicate that a signal is present at the receiver inputs.

5.11.1.3.2 Semantics

The SIGNAL.indication primitive provides parameters as follows:

SIGNAL.indication (Signal Status).

5.11.1.3.3 When Generated

The SIGNAL indication primitive shall be passed to the lane layer when a the receiver detects that the status of the received signal changes.

5.11.1.3.4 Effect On Receipt

The effect on receipt of the SIGNAL.indication primitive by the lane layer shall be to notify the Lane Initialisation state machine that of the status of the signal at the receiver.

NOTE SIGNAL.indication is used to provide the LossOfSignal signal used in the Lane Initialisation and Standby Management state machine.

5.11.2 8B/10B encode/decode

a. The SpaceFibre CODEC shall use 8B/10B encoding to encode each 8-bit data character or control code into a 10-bit symbol that is transmitted.

NOTE 8B/10B encoding is described in A2.2.

b. To ensure DC balancing of the transmitted signal account shall be kept of the current running disparity in the transmitter.

c. If the current running disparity is positive when encoding an 8-bit character or control code, the symbol for that data character or control code which has negative disparity shall be used.

d. If the current running disparity is positive when encoding an 8-bit character or control code, and there is no symbol for that data character or control code which has negative disparity, the symbol with neutral shall be used.

e. If the current running disparity is negative when encoding an 8-bit character or control code, the symbol for that data character or control code which has positive disparity shall be used.

f. If the current running disparity is negative when encoding an 8-bit character or control code, and there is no symbol for that data character or control code which has positive disparity, the symbol with neutral shall be used.

g. To detect disparity errors account shall be kept of the current running disparity in the receiver.

h. If the current running disparity is more than plus one or less minus one, this shall indicate a disparity error.

i. If a disparity error occurs, it shall be indicated to the receive synchronisation state machine and the current symbol shall be set to K0.0.

j. When a symbol is received it shall be decoded into an 8-bit data character or control code using the 8B/10B symbol table.

k. If an unrecognised symbol is received then a symbol error shall be indicated to the receive synchronisation state machine and the current symbol shall be set to K0.0.

l. The 8B/10B encoding shall with the least-significant five bits being encoded as detailed in Table 5-30 and the most significant three bits being encoded as detailed in Table 5-31.

Table 5-30 5B/6B Encoding			
Input		**Output**	
Data Input	**Data bits 43210 (EDCBA)**	**Current Running Disparity -ve abcdei**	**Current Running Disparity +ve abcdei**
D00.y	00000	100111	011000
D01.y	00001	011101	100010
D02.y	00010	101101	010010
D03.y	00011	110001	
D04.y	00100	110101	001010
D05.y	00101	101001	
D06.y	00110	011001	
D07.y	00111	111000	000111
D08.y	01000	111001	000110
D09.y	01001	100101	
D10.y	01010	010101	
D11.y	01011	110100	
D12.y	01100	001101	
D13.y	01101	101100	
D14.y	01110	011100	
D15.y	01111	010111	101000
D16.y	10000	011011	100100
D17.y	10001	100011	
D18.y	10010	010011	
D19.y	10011	110010	
D20.y	10100	001011	
D21.y	10101	101010	
D22.y	10110	011010	
D/K23.y	10111	111010	000101
D24.y	11000	110011	001100
D25.y	11001	100110	
D26.y	11010	010110	
D/K27.y	11011	110110	001001
D28.y	11100	001110	
K28.y	11100	001111	110000
D/K29.y	11101	101110	010001
D/K30.y	11110	011110	100001
D31.y	11111	101011	010100

Table 5-31 3B/4B Encoding			
Input		Output	
Data Input	Data bits 765 (HGF)	5B/6B Disparity -ve fghj	5B/6B Disparity +ve fghj
D/Kxx.0	000	1011	0100
Dxx.1	001	1001	
Kxx.1	001	0110	1001
Dxx.2	010	0101	
Kxx.2	010	1010	0101
D/Kxx.3	011	1100	0011
D/Kxx.4	100	1101	0010
Dxx.5	101	1010	
Kxx.5	101	0101	1010
Dxx.6	110	0110	
Kxx.6	110	1001	0110
Dxx.7	111	1110/0111	0001/1000
Kxx.7	111	0111	1000

5.11.3 Symbol synchronisation

a. The boundary between symbols shall be determined by detecting the unique comma sequences.

b. The **Positive Comma** sequence is 0011111.

c. The **Negative Comma** sequence is 1100000.

NOTE The full 7 bit comma sequences are used for symbol synchronisation.

d. Both positive and negative commas shall be detected and used for synchronisation.

e. Synchronisation may be performed on only positive commas or only negative commas when using legacy SerDes devices.

f. The 10-bit wide input stream shall be aligned or realigned to form a stream of correctly aligned symbols so that each 10-bit group contains one complete symbol.

g. The 10-bit wide input stream shall be realigned every time a comma sequence is detected in a different position to the position of the last comma detected.

h. When a 20-bit or 40-bit interface is being used from the 8B/10B receiver, realignment may occur on the first comma in a word, i.e. when there are two or more commas in a 20-bit or 40-bit word it will be the comma in the lower significant bit position that is used for symbol realignment.

5.11.3.2 Receive Synchronisation State Machine

a. A receive synchronisation state machine shall be used to determine when incoming symbols are properly synchronised.

NOTE The state diagram for the receive synchronisation state machine is illustrated in Figure 5-9.

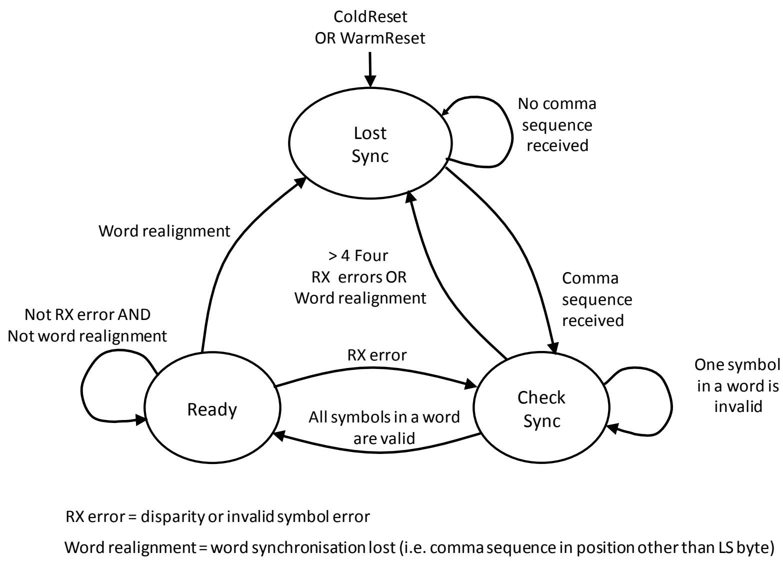

Figure 5-9 Receive Synchronisation State Machine

5.11.3.2.1 LostSync State

a. The LostSync state shall be entered on one of the following conditions:

1. ColdReset command.
2. WarmReset command.
3. From the CheckSync state, when word realignment occurs or when a total of more than four symbols are received that are invalid or contain a disparity error.
4. From the Ready state, when word realignment occurs.

b. When in the LostSync state the Receive Synchronisation state machine shall initiate the following actions:

1. Replace any received data and control words with RXERR control words.

c. The Receive Synchronisation state machine shall leave the LostSync State on one of the following conditions, which shall be evaluated in the order given:

   1. When a comma sequence is received, move to the CheckSync state.

d. The LostSync state is summarised in Table 5-32.

Table 5-32 LostSync State

State	LostSync
Entry	ColdReset command. WarmReset command. From the CheckSync state, when word realignment occurs or when a total of more than four symbols are received that are invalid or contain a disparity error. From the Ready state, when word realignment occurs.
Action	Replace received data and control words with RXERR control words.
Exit	When a comma sequence is received, move to the CheckSync state.

5.11.3.2.2 CheckSync State

a. The CheckSync state shall be entered on one of the following conditions:

   1. From the LostSync state, when a comma sequence is received.

   2. From the Ready state, when a disparity error occurs or an invalid symbol is detected.

b. When in the CheckSync state the Receive Synchronisation state machine shall initiate the following actions:

   1. Replace any received words with RXERR control words.

   2. Count the number of symbols received that are invalid or contain a disparity error.

c. The Receive Synchronisation state machine shall leave the CheckSync State on one of the following conditions, which shall be evaluated in the order given:

   1. When ColdReset is asserted, move to the LostSync state.

   2. When WarmReset is asserted, move to the CheckSync state.

   3. When word realignment occurs, move to the LostSync state.

   4. When a total of more than four symbols are received that are invalid or contain a disparity error, move to the LostSync state.

      NOTE These four invalid or erroneous symbols need not be sequential i.e. there may be valid symbols interspersed amongst them

5. When all symbols within a word are valid, move to the Ready state.

d. The CheckSync state is summarised in Table 5-33.

Table 5-33 CheckSync State

State	CheckSync
Entry	From the LostSync state, when a comma sequence is received. From the Ready state, when a disparity error occurs or an invalid symbol is detected.
Action	Replace any received words with RXERR control words. Count the number of symbols received that are invalid or contain a disparity error.
Exit	When ColdReset is asserted, move to the LostSync state. When WarmReset is asserted, move to the CheckSync state. When word realignment occurs, move to the LostSync state. When a total of more than four symbols are received that are invalid or contain a disparity error, move to the LostSync state. When all symbols within a word are valid, move to the Ready state.

5.11.3.2.3 Ready State

a. The Ready state shall be entered on one of the following conditions:

   1. From the CheckSync state, when all symbols within a word are valid.

b. When in the Ready state the Receiver Synchronisation state machine shall initiate the following actions:

   1. Receive symbols.

c. The Receiver Synchronisation state machine shall leave the Ready state on one of the following conditions, which shall be evaluated in the order given:

   1. When ColdReset is asserted, move to the LostSync state.
   2. When WarmReset is asserted, move to the CheckSync state.
   3. When a disparity error occurs or an invalid symbol is received, move to the CheckSync state.
   4. When word realignment occurs, move to the LostSync state.

d. The Ready state is summarised in Table 5-34.

Table 5-34 Ready State

State	Ready
Entry	From the CheckSync state, all symbols within a word are valid.
Action	Receive symbols.

Exit	When ColdReset is asserted, move to the LostSync state. When WarmReset is asserted, move to the CheckSync state. When a disparity error occurs or an invalid symbol is received, move to the CheckSync state. When word realignment occurs, move to the LostSync state.

5.11.4 Word Synchronisation

a. A symbol word shall comprise four symbols.

b. The first symbol to be transmitted in a symbol word shall be the least significant symbol.

c. A comma shall only occur in the least significant symbol position of a symbol word.

d. Word synchronisation shall be performed whenever a comma sequence is received.

e. Word synchronisation shall be achieved by selecting symbols in consecutive groups of four so that the comma sequence appears in the least significant symbol position.

f. On word synchronisation, the first word shall comprise the comma control symbol in the least significant symbol position, together with the following three symbols.

g. Each subsequent group of four symbols shall form each of the following symbol words.

h. Word realignment shall occur when a comma sequence occurs in any position other than the least significant symbol position of a word.

i. If a word contains a K0.0 symbol indicating an error that word together with the previous word shall be set to the RXERR control word.

NOTE An 8B10B disparity error will be detected in the next symbol that has disparity of +/- 1. All control words have disparity except EDF which has zero disparity. This ensures that a data frames, containing an error will not be passed to the upper layers without an error being flagged.

5.12 Serialisation layer

5.12.1 SerDes interface

a. The SerDes Interface shall pass coded, but unsynchronised, symbols between the Encoding and Serialisation parts of the SpaceFibre CODEC.

b. The SerDes interface shall comprise a Transmit SerDes interface and a Receive SerDes interface.

c. The SerDes interfaces shall be 10-bit, 20-bit or 40-bit wide.

5.12.1.2 Transmit SerDes Input

The Transmit SerDes Input shall contain the signals listed in Table 5-35, Table 5-36, or Table 5-37.

Table 5-35 Transmit SerDes Interface (10-bit)

Signal	Dir	Function
SerDes_Txdata(9:0)	In	10-bit wide data containing one symbol for transmission.

Table 5-36 Transmit SerDes Interface (20-bit)

Signal	Dir	Function
SerDes_Txdata(19:0)	In	20-bit wide data containing two symbols for transmission.

Table 5-37 Transmit SerDes Interface (40-bit)

Signal	Dir	Function
SerDes_Txdata(39:0)	In	40-bit wide data containing four symbols for transmission.

5.12.1.3 Receive SerDes Output

The Receive SerDes Output shall contain the signals listed in Table 5-38, Table 5-39, or Table 5-40.

Table 5-38 Receive SerDes Interface (10-bit)

Signal	Dir	Function
SerDes_Rxdata(9:0)	Out	10-bits of received data. This data is NOT symbol synchronised i.e. the 10-bits can contain some bits from one symbol and the rest of the bits from the next symbol.
RX_Signal	OUT	1 bit that indicates if there is a signal at the receiver.

Table 5-39 Receive SerDes Interface (20-bit)

Signal	Dir	Function

SerDes_Rxdata(19:0)	Out	20-bits of received data. This data is NOT symbol synchronised i.e. the 20-bits can contain some bits from one symbol, the following complete symbol, and the rest of the bits from the next symbol.
RX_Signal	OUT	1 bit that indicates if there is a signal at the receiver.

Table 5-40 Receive SerDes Interface (40-bit)

Signal	Dir	Function
SerDes_Rxdata(39:0)	Out	40-bits of received data. This data is NOT symbol synchronised i.e. the 40-bits can contain some bits from one symbol, the following complete three symbols, and the rest of the bits from the next symbol.
RX_Signal	OUT	1 bit that indicates if there is a signal at the receiver.

5.12.2 Bit Synchronisation

a. The receive clock used to sample the incoming bit stream, shall be generated by a clock recovery circuit that matches the phase of a local receive clock to the transitions of the incoming bit stream.

b. Sampling of the bit stream should be close (+/- 20% TBC) to the centre of the bit period.

c. The clock recovery circuit should indicate in a status register when bit synchronisation is achieved.

NOTE This is for status information only.

5.12.3 Serialiser/deserialiser

a. 10-bit symbols shall be transmitted serially over the physical medium.

b. At the transmitter a serialiser shall be used to convert each parallel 10-bit symbol into a serial bit stream with each bit of the symbol being send one after the other.

NOTE One, two or four symbols can be provided to the serialiser in parallel.

c. The least significant bit of the 10-bit symbol shall be transmitted first.

d. At the receiver the incoming bit stream shall be converted to a 10-bit word by sampling the bit stream with a receive clock in a deserialiser.

NOTE The deserialiser does not necessarily produce a stream of 10-bit symbols because the boundary of the symbols is not known by the deserialiser.

5.12.4 Inversion

The received symbols shall be bit-wise inverted if requested by the lane initialisation state machine.

5.12.5 Serial loopback

a. A serial loopback facility shall be provided in the SpaceFibre CODEC for test purposes.

b. The serial loopback shall connect the bit stream output from the serialiser in the transmitter directly into the serial input of the deserialiser in the receiver.

c. The serial loopback shall connect the bit stream input from the receiver directly to the transmitter driver output.

5.13 Physical layer

5.13.1 Serial interface

a. The Serial Interface shall pass serial data out of and into the SpaceFibre CODEC.

b. The serial interface shall comprise a transmitter serial output and a receiver serial input.

5.13.1.2 Transmit serial output

The Transmit Serial Output shall contain the signals listed in Table 5-41.

Table 5-41 Transmit Serial Interface

Signal	Function
Txp	Positive side of the differential serial transmitter output.
Txn	Negative side of the differential serial transmitter output.

5.13.1.3 Receive serial input

The Receive Serial Input shall contain the signals listed in Table 5-42.

Table 5-42 Receiver Serial Interface

Signal	Function
Rxp	Positive side of the differential serial receiver input.
Rxn	Negative side of the differential serial receiver input.

5.13.2 Electrical SpaceFibre medium

5.13.2.1 Electrical SpaceFibre driver and receiver

a. The driver and receiver for SpaceFibre operation over copper shall use Current Mode Logic (CML), or compatible driver and receiver.

b. The eye pattern at the outputs of the transmitter shall be as illustrated in Figure 5-10.

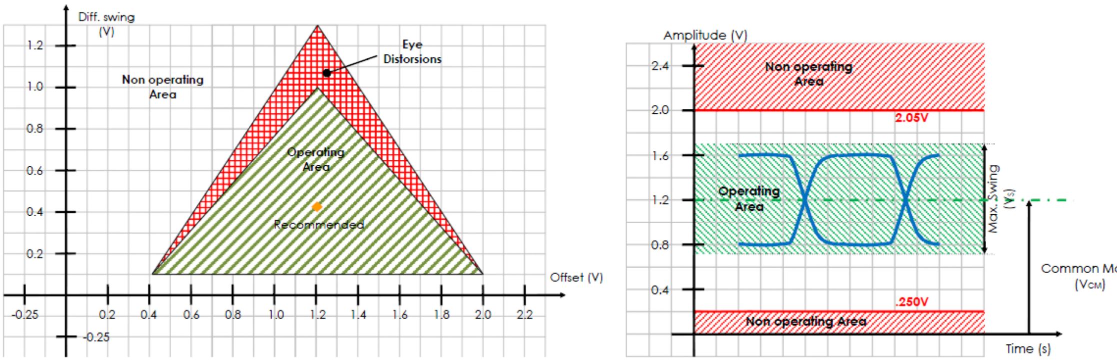

(a) Diff. Vs Common mode voltage (b)Time representation

Eye pattern to be provided.

Figure 5-10 Transmitter Eye Pattern

c. The eye pattern at the inputs to the receiver shall be as illustrated in Figure 5-11.

Eye pattern to be provided.

Figure 5-11 Receiver Eye Pattern

5.13.2.2 Electrical SpaceFibre PCB tracks

a. The PCB tracks for electrical SpaceFibre shall be 100 ohms differential impedance +/- 5 ohms.

b. Two pairs of differential PCB tracks shall be used for a bi-directional SpaceFibre link, one pair for each direction.

5.13.2.3 EGSE electrical connectors

a. For electrical ground support equipment (EGSE) external serial ATA connectors shall be used as specified in Serial ATA Revision 3.0, clause 6.5.1.

5.13.2.4 EGSE cable

a. For electrical ground support equipment (EGSE) external serial ATA cable shall be used specified in Serial ATA Revision 3.0, clause 6.6.1.

5.13.2.5 EGSE cable assemblies

a. For electrical ground support equipment (EGSE) a crossover external serial ATA cable shall be used.

b. The EGSE cable assembly shall be terminated at each end by an External Serial ATA cable receptacle as detailed in Serial ATA Revision 3.0, clause 6.5.1 and Figure 93 of that specification.

c. The EGSE cable assembly shall use the external serial ATA cable specified clause 5.13.2.4.

d. The EGSE cable assembly shall be wired as illustrated in Figure 5-12.

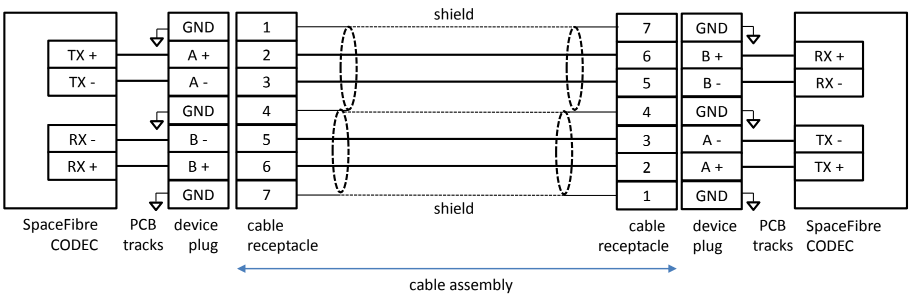

Figure 5-12 SpaceFibre EGSE cable assembly

5.13.2.6 Flight electrical connectors

a. For flight equipment connections conforming to ESCC specification (draft 07072-ST-MDSA HDR -01) shall be used.

5.13.2.6.1 Male connectors

a. Male two way connectors (variant 05) shall be used for cable assemblies.

5.13.2.6.2 Female connectors

a. Female two way connectors (variant 08) shall be used for EGSE to flight adaptor cable assemblies.

5.13.2.6.3 PCB mounting female connectors

a. Female two way PCB panel mounting connectors (variant 02) or female two way PCB mounting connectors (variant 11) shall be used for connecting to a PCB.

5.13.2.7 Flight cable

a. Flight cable shall be PTFE coaxial cable that conforms to the cable specification in ESCC specification (draft 07072-ST-MDSA HDR -01), clause 4.4.7.

5.13.2.8 Flight cable assemblies

a. Male connectors 2 way connectors conforming to clause 5.13.2.6 shall be used on flight cable assemblies.

b. The flight cable assembly shall use the co-axial cable specified in clause 5.13.2.7.

c. The flight cable assembly shall be wired as illustrated in Figure 5-13.

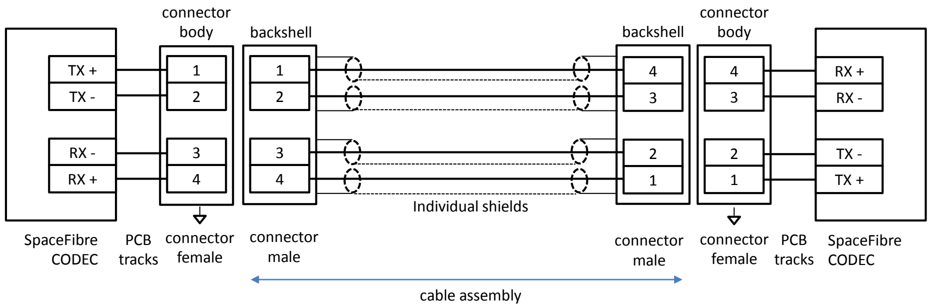

Figure 5-13 SpaceFibre flight cable assembly

5.13.2.1 EGSE to flight adaptor cable assemblies

a. The EGSE to flight adaptor cable assembly shall allow a SpaceFibre EGSE unit to connect to a flight cable assembly.

b. One end of the EGSE to flight adaptor cable assembly shall be terminated by an External Serial ATA cable receptacle as detailed in Serial ATA Revision 3.0, clause 6.5.1 and Figure 93 of that specification.

c. The other end of the EGSE to flight adaptor cable assembly shall be terminated by a female 2 way connector conforming to clause 5.13.2.6.2.

d. The EGSE to flight adaptor cable assembly shall use the co-axial cable specified in clause 5.13.2.7.

e. The EGSE to flight adaptor cable assembly shall be wired as illustrated in Figure 5-14.

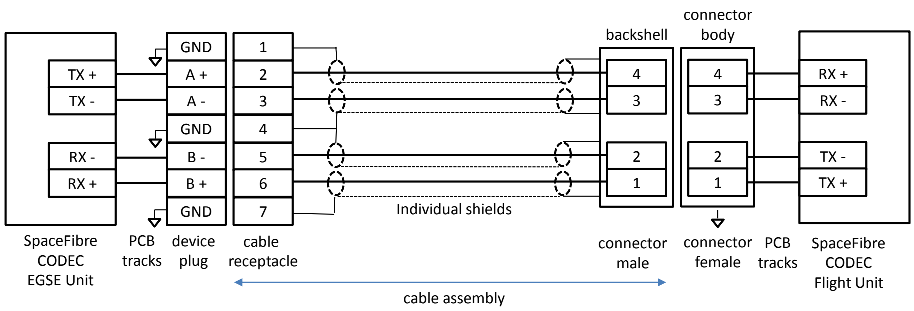

Figure 5-14 SpaceFibre EGSE to flight adaptor cable assembly

f. The lengths of the co-axial cables used in the cable assembly shall be matched to within 1% and to within +/- 3 mm maximum.

5.13.2.2 Flight connector savers

a. One male 2 way connector and one female 2 way connector conforming to clause 5.13.2.6 shall be used on flight connector savers.

b. The flight connector shall use the co-axial cable specified in clause 5.13.2.7 to wire between the connectors.

c. The flight connector saver shall be wired as illustrated in Figure 5-13.

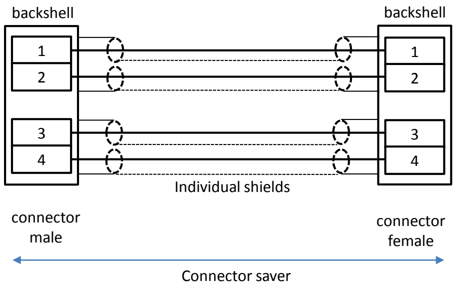

Figure 5-15 SpaceFibre flight connector saver

d. The length of the co-axial cables used in the flight connector shall be of the same length within +/- 1 mm and of maximum length 100 mm.

5.13.3 Fibre optic driver and receiver

5.13.3.1 Fibre optic driver and receiver

a. TBA

5.13.3.2 Fibre optic connectors

a. TBA

5.13.3.3 Fibre optic cables

a. TBA

5.14 Management layer

a. The management layer shall provide a means of configuring the SpaceFibre CODEC and reading its status.

b. Configuration of the SpaceFibre CODEC shall be performed by writing to management parameters via the management interface.

c. Reading of the status of the SpaceFibre CODEC shall be performed by reading status values via the management interface.

5.14.2 Configuration parameters

a. It shall be possible to read and write configuration parameters via the management interface.

b. The configuration parameters are listed in Table 5-43.

Table 5-43 SpaceFibre CODEC Configuration Parameters

Layer	Parameter	Description	Default value
Virtual Channel	VC number (VC buffer)	Virtual Channel number assigned to the hardware buffer identified by the VC buffer parameter. This parameter shall only be configured following a cold reset.	0
	Priority level (VC buffer)	Priority level (0-15) assigned to a particular virtual channel.	15
	Expected Bandwidth (VC buffer)	The fraction of overall link bandwidth assigned to a particular virtual channel.	1
	Number of Time Slots	Number of time slots in the schedule. The maximum number of time slots in a schedule shall be 256.	64
	Allocated Time Slots (VC buffer)	The time slots in which a particular virtual channel is permitted to send data frames. This parameter is an array of N bits, where N is the number of time slots in the schedule. A bit is set to one to indicate that the virtual channel can send data in the corresponding time slot.	All ones
	Virtual Channel Idle Time Limit	Optional parameter. Determines the maximum time a channel can be idle without sending data before the underuse indication is raised. The period value can be hardcoded or can be provided by an optional configuration parameter for all VC channels or for each VC channel independently.	1 ms
Multi-Lane layer	Required lanes	Required number of lanes to be used to form the SpaceFibre link.	
Lane Layer	Start mode (lane)	Asserts or deasserts Lane_Start for the corresponding lane.	0
	Autostart (lane)	Asserts or deasserts AutoStart for the corresponding lane.	0

	RXERR counter decrement period (lane)	OPTIONAL parameter. When not provided the value should be hardcoded. Period of time (in number of words received) before the RXERR words counter is decremented by one. It determines the minimum BER. When the error rate of the lane is higher than provided the lane will be reset and restarted.	1024

5.14.3 Status parameters

a. It shall be possible to read and write configuration parameters via the management interface.

b. The mandatory status parameters shall be as listed in Table 5-44.

c. Implementations may provide additional status information for debugging purpose.

NOTE Basic protocol and CODEC debugging can be done using just the mandatory status parameters.

d. The way the status parameters are cleared shall be implementation dependent.

Table 5-44 SpaceFibre CODEC Status Parameters		
Layer	Parameter	Description
Virtual Channel	BW over use (VC buffer)	Indicates that the hardware buffer is using much more bandwidth than expected and has reached the minimum bandwidth credit value.
	BW under use (VC buffer)	Indicates that the hardware buffer is using less bandwidth than expected and that the bandwidth credit has remained at its maximum value for a certain time (the maximum expected Idle time).
	Has Credit (VC buffer)	Indicates that there is space in the input buffer at the other end of the link (destination).
	Input buffer overflow (VC buffer)	Indicates that an input buffer is receiving data when it is full. This should not happen and indicates a fatal protocol error.
Broadcast layer	BC missed	Set when the broadcast sequence is two more than the last value.
	Sequence error	Set when the broadcast sequence is different and not one or two more than the last value received. Should not be set when it is the first value received after a cold reset.
Retry Layer	CRC-16 error	Set when a CRC-16 error has occurred. This indicates that one or multiple lanes are producing multiple bit errors, which is not expected under nominal operation.
	CRC-8 or sequence error	Set when a CRC-8 or a sequence error occurs. This can occur under nominal operation.
	Retry buffer empty	Set when the retry buffer is empty. This indicates that all data from the virtual channels has been sent and acknowledged.
	Number of retries	The number of retries made by the SpaceFibre CODEC. This is cleared on cold reset and remote flush and incremented by one for every retry event initiated.
Lane Layer	Lane Active (lane)	Set when the lane is in active state
	LossOfSignal (lane)	Set when LossOfSignal is detected in Active state.
	RxError	Set when one RxError occurs in Active state.

	(lane)	This can occur under nominal operation.
	RxError overflow (lane)	Set when RxError counter overflows in Active state.
	Standby (lane)	Set when Standby words are received.
	Timeout (lane)	Set when a connection timeout occurs.
	LOS (lane)	Set when LOS words are received.
	Remote Flush (lane)	Set when a Remote Flush is received.
	RX Polarity (lane)	Set when the receiver polarity is inverted.
Serialisation	Bit Sync	Set when bit synchronisation is achieved.

5.14.4 Reset

a. The effects of cold reset on the SpaceFibre CODEC shall be as indicated in Table 5-45.

b. The effects of warm reset on the SpaceFibre CODEC shall be as indicated in Table 5-45.

Table 5-45 Effects of Cold Reset and Warm Reset

Layer	Variable	Cold Reset	Warm Reset
Virtual Channel	Output VCBs	Flushed	Unchanged
	Input VCBs	Flushed	Unchanged
	FCT counter	Cleared	Unchanged
	Bandwidth credit	Set to zero	Unchanged
	Input space counter	Set to buffer size	Unchanged
	FCT credit counter	Set to zero	Unchanged
Broadcast Message	BC Sequence counter	Set to zero	Unchanged
Framing			
Retry	Transmit frame sequence counter	Set to zero	Unchanged
	Receive frame sequence counter	Set to zero	Unchanged
	Number of retries	Set to zero	Set to zero
	Retry buffers	Flushed	Unchanged
	PRBS seed	Set to 0xffff ffff	Unchanged
Lane Control	Required number of lanes	Set to one	Unchanged
	Distribution lane numbers	Clear	Clear
	Concentration lane numbers	Clear	Clear
Lane	Capability parameters	Reset	Unchanged
	Receiver bit inversion	Off	Off
Encoding			
Serialisation	PLL	Reset	Reset
Physical			

5.15 SpaceFibre conformance

5.15.1 Overview

5.15.2 Partial implementations

Single lane

Bit inversion

Only positive or only negative comma synchronisation

Parallel loopback

Annex A(informative) Serial Data Link Concepts

This section provides an overview of several key concepts for high-speed serial data links.

A.1 Data Scrambling

Data scrambling is a technique used to reduce the electro-magnetic (EM) emissions from a communications system. The data signal is convolved with a wideband signal which results in the spectrum of the data being broadened. Possible peaks in the EM spectrum of the original data signal are spread out reducing the energy at any single frequency. Note that data scrambling does not guarantee reduced peaks in the EM spectrum since it is possible that the scrambling produces a bit sequence with a higher spectral peaks than the original signal. However, for regular bit sequences it is likely to reduce the spectral peaks.

A random number generator is used to produce the wideband signal which is XORed with the data being transmitted. At the beginning of each frame being transmitted the random number generator is reseeded with a specific value. A similar random number generator in the receiver, seeded with the same seed as in the transmitter at the start of every new frame, is used to de-scramble the data. The incoming data is XORed with the random number sequence to reveal the original data stream.

The random number generator is implemented using a linear feedback shift register as shown in Figure A-1. An example scrambling/de-scrambling polynomial is that used in PCI-Express:

$$G(x) = X^{16} + X^{5} + X^{4} + X^{3} + 1$$

The seed for the random number generator is $FFFF_h$ *i.e.* all flip-flops in the random number generator are set to 1.

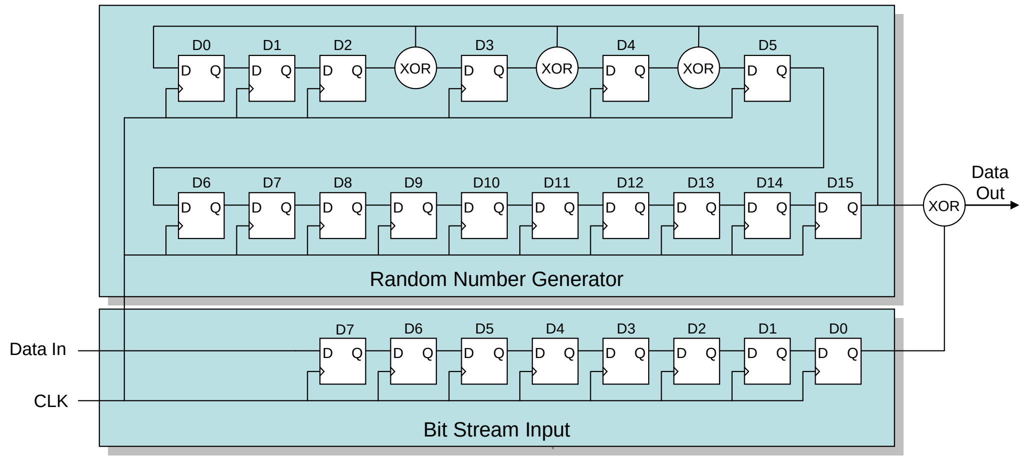

Figure A-1 Scrambler / De-Scrambler

A.2 8B/10B Encoding and Decoding

8B/10B encoding encodes 8-bit data bytes into 10-bit characters for transmission. The 8B/10B encoding has several advantages over direct 8-bit transmission.

1. It provides a transmitted data stream with roughly the same number of 1's as 0's giving the data a zero DC bias, improving the transmission characteristics and enabling AC coupling.
2. Since a 10-bit code has 1024 possible values and not all of these are needed to send an 8-bit value there are spare valid codes left over that can be used for control codes.
3. It guarantees that there will be sufficient number of bit transitions in the serial data stream to enable the recovery of the bit clock using a phase-locked loop. A maximum of five consecutive ones or zeros are ensured with 8B/10B encoding.
4. Since all characters, both data and control characters, are transmitted with 10-bits the bit and character transmission rates are constant simplifying the transmission and reception of characters.
5. Codes that are unused by the 8B/10B encoding can be used to detect link errors *i.e.* if an unused code occurs then there has been a transmission error.
6. The current running disparity following 8B/10B encoding is always +1 or -1, any other value indicates a disparity error.

To avoid significant DC components 8B/10B encoding uses only the 10-bit codes that contain either 5 ones and 5 zeros, 6 ones and 4 zeros, or 4 ones and 6 zeros. There are enough of these to encode the 8-bit data byte and several possible control codes. Characters encoded with 5 ones and 5 zeros have neutral disparity and will produce zero DC bias. However, if a sequence of bytes was transmitted that contained characters all with 6 ones and 4 zeros the DC component would slowly increase. A similar opposite effect would occur if the characters all contained 4 ones and 6 zeros. To prevent this increasing DC bias and to maintain an equal number of transmitted ones and zeros each character with an unequal number of ones and zeros has two possible codes one with 6 ones and 4 zeros and the other with 4 ones and 6 zeros.

Every time the transmitter sends a character with 6 ones and 4 zeros it will record the fact that it has sent more ones than zeros and the next time it has to send a character with an uneven number of bits it will choose the code that has 4 ones and 6 zeros. This keeps the average number of ones and zeros the same and eliminates any DC bias in the transmitted signal. The Current Running Disparity variable is set to one (positive) when more ones have been sent than zeros and to zero (negative) when more zeros have been sent than ones. Characters with 5 ones and 5 zeros have neutral disparity and do not affect the Current Running Disparity value. When a character with an unequal number of ones and zeros is to be sent, the value of the Current Running Disparity will determine which of the two possible 10-bit codes will be sent. If the Current Running Disparity is positive then the option with 4 ones and 6 zeros is sent, if it is negative then the other option with 6 ones and 4 zeros is transmitted.

Once all 256 possible values of an 8-bit data byte have been assigned a code with 5 ones and 5 zeros or a pair of codes with unequal numbers of ones and zeros, there are just 12 valid codes left out of the possible 1024 values of a 10-bit code. The others have more than six ones or more than six zeros and are invalid.

A.2.1 8B/10B Encoding

8B/10B encoding is normally done using a pair of look-up tables as shown in Figure A-2 rather than a single look-up table.

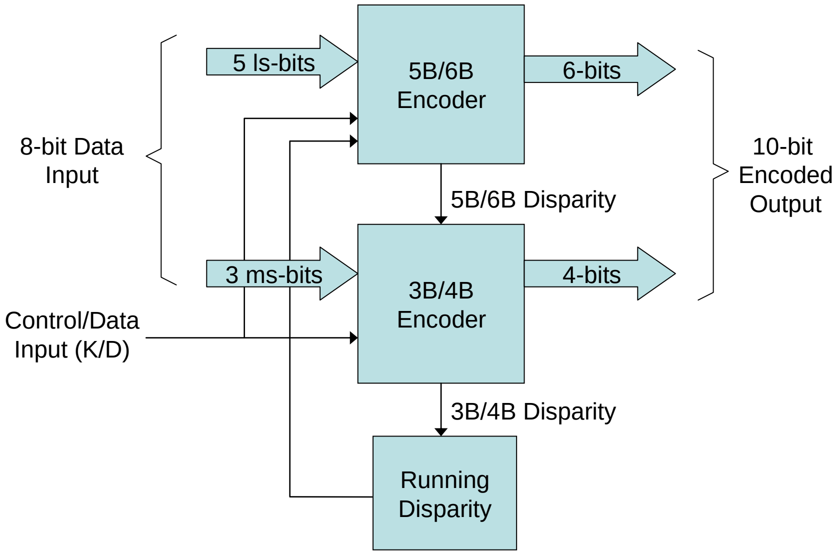

Figure A-2 8B/10B Encoder

The 5B/6B and 3B/4B approach to 8B/10B encoding has lead to a specific notation for representing codes resulting from this encoding. This is illustrated in Figure A-3.

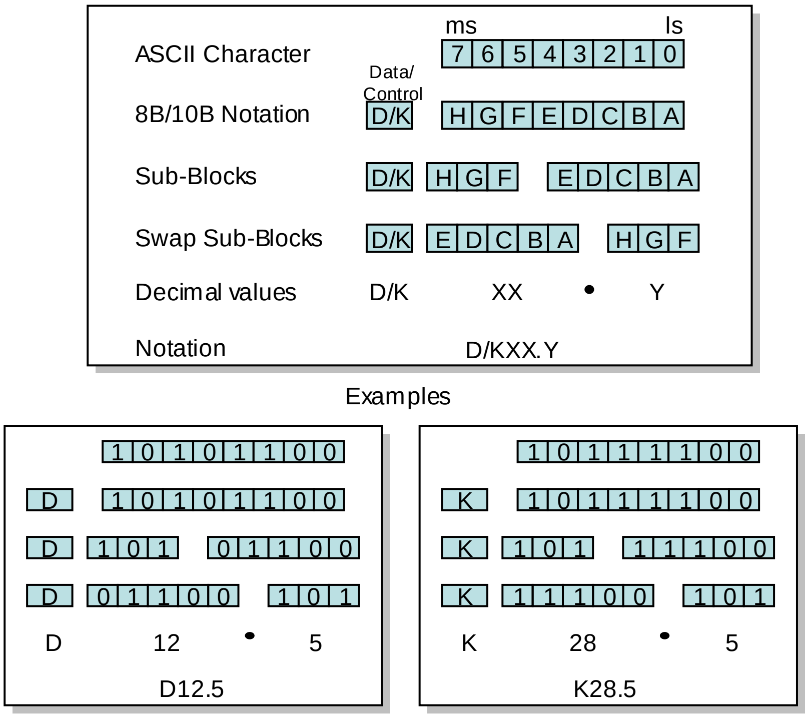

Figure A-3 8B/10B Notation

The five least-significant bits are encoded first using a 5B/6B encoder. This takes into account whether the 5 least significant bits are part of a control or data word as determined by the K/D input and also the current running disparity of the link. The 5B/6B encoding table is given in Table A-1. This table has the following properties:

- The six bit outputs consist of either three ones and three zeros, four ones and two zeros or two ones and four zeros.
- When the output code has neutral disparity (three ones and three zeros) there is one code independent of the running disparity (except for D07.y which has two codes based on the running disparity). The complement of a neutral disparity code also has neutral disparity and, except for D07.y, corresponds to a different input symbol. By definition, using a neutral disparity code will not affect the running disparity.
- When the output code has non-neutral disparity (four ones and two zeros or two ones and four zeros) there are two alternative codes provided which are the complement of each other. The code that is applied when the current running disparity of the link is negative (-ve) has four ones and two zeros which will then make the disparity out of the 5B/6B encoder positive. Similarly when the current running disparity is positive (+ve) the code with two ones and four zeros is applied making the 5B/6B disparity negative.
- The coding table is organised to minimise the amount of logic needed to implement it so that wherever possible there is a one to one mapping of bits from the 5-bit input to the 6-bit output. Note that K28.y must be treated as a special case.

The six-bit output of the 5B/6B encoder forms the six least-significant bits of the 8B/10B encoder output. The 5B/6B disparity is used in the encoding of the three most-significant bits of the 8-bit input data. The 5B/6B disparity, three most significant bits of the input data and the control/data flag (K/D) are fed into a separate 3B/4B encoder which produces the four most significant bits of the 8B/10B encoder output and a new value for the running disparity. The contents of the 3B/4B encoding table are given in Table A-2. This table has the following properties:

- Only 13 possible codes are valid, those shown in Table A-2.
- When the output code has non-neutral disparity there are two codes which are the complement of each other (with the exception of the codes for Dxx.7). One of these two codes will be used depending on the 5B/6B disparity. If the 5B/6B disparity is negative then the option with three ones and one zero will be used resulting in an overall positive disparity which will be the new value of the running disparity. The opposite is the case when the 5B/6B disparity is positive.
- The encoding for Dxx.7 has an alternative coding to prevent five consecutive ones being transmitted. The -ve current running disparity alternative (0111_b) is used for D17.7, D18.7 and D20.7. The +ve current running disparity alternative (1000_b) is used for D11.7, D13.7 and D14.7. This does complicate the encoding somewhat because these special cases have to be identified in the input data stream and the alternative code activated.

The complete 8B/10B encoding is performed by combining the results of the 5B/6B and 3B/4B encoding steps.

A.2.2 8B/10B Decoding

The task of decoding 8B/10B symbols is more complicated than the encoding process since a large number of input codes are mapped to a few valid output codes. Care must be taken to ensure that invalid 8B/10B codes are not accidentally considered to be valid simply because the 5B/6B and 3B/4B components are individually valid. An example is 1110101110_b which has a valid 5B/6B component 111010_b (-D23.y) and a valid 3B/4B component 1110_b (Dxx.7 normal encoding). This is invalid because the alternative Dxx.7 encoding 0111_b encoding ought to have been used.

Additional care must be taken with the 3B/4B decoder since the 4-bit input cannot distinguish between K and D codes. For example, 0110_b represents either -Kxx.1, or Dxx.6 or +Kxx.6.

Table A-1 5B/6B Encoding			
Input		Output	
Data Input	Data bits 43210 (EDCBA)	Current Running Disparity -ve abcdei	Current Running Disparity +ve abcdei
D00.y	00000	100111	011000
D01.y	00001	011101	100010
D02.y	00010	101101	010010
D03.y	00011	110001	
D04.y	00100	110101	001010
D05.y	00101	101001	
D06.y	00110	011001	
D07.y	00111	111000	000111
D08.y	01000	111001	000110
D09.y	01001	100101	
D10.y	01010	010101	
D11.y	01011	110100	
D12.y	01100	001101	
D13.y	01101	101100	
D14.y	01110	011100	
D15.y	01111	010111	101000
D16.y	10000	011011	100100
D17.y	10001	100011	
D18.y	10010	010011	
D19.y	10011	110010	
D20.y	10100	001011	
D21.y	10101	101010	
D22.y	10110	011010	
D/K23.y	10111	111010	000101
D24.y	11000	110011	001100
D25.y	11001	100110	
D26.y	11010	010110	
D/K27.y	11011	110110	001001
D28.y	11100	001110	
K28.y	11100	001111	110000
D/K29.y	11101	101110	010001
D/K30.y	11110	011110	100001
D31.y	11111	101011	010100

Table A-2 3B/4B Encoding			
Input		**Output**	
Data Input	**Data bits 765 (HGF)**	**5B/6B Disparity -ve fghj**	**5B/6B Disparity +ve fghj**
D/Kxx.0	000	1011	0100
Dxx.1	001	1001	
Kxx.1	001	0110	1001
Dxx.2	010	0101	
Kxx.2	010	1010	0101
D/Kxx.3	011	1100	0011
D/Kxx.4	100	1101	0010
Dxx.5	101	1010	
Kxx.5	101	0101	1010
Dxx.6	110	0110	
Kxx.6	110	1001	0110
Dxx.7	111	1110/0111	0001/1000
Kxx.7	111	0111	1000

The 12 control characters are listed in Table A-3. Three of these characters (K28.1, K28.5 and K28.7) contain a unique seven bit pattern (0011111 or 1100000) which does not occur in any of the data codes and which cannot be produced by concatenating any other two data or control codes. This pattern is known as the "comma" pattern and is widely used for performing receive code synchronisation (character alignment). The comma pattern is underlined in Table A-3.

Note that K28.7 followed by certain other data or control codes can produce a false comma, but the correct one comes first.

Table A-3 8B/10B Control (K) Codes		
Input	**Output**	
Special Character Name	**Current Running Disparity -ve**	**Current Running Disparity +ve**
K28.0	001111 0100	110000 1011
K28.1	001111 1001	110000 0110
K28.2	001111 0101	110000 1010
K28.3	001111 0011	110000 1100
K28.4	001111 0010	110000 1101
K28.5	001111 1010	110000 0101
K28.6	001111 0110	110000 1001
K28.7	001111 1000	110000 0111
K23.7	111010 1000	000101 0111
K27.7	110110 1000	001001 0111
K29.7	101110 1000	010001 0111
K30.7	011110 1000	100001 0111

A.2.3 Disparity

The initial disparity can be either positive or negative, i.e. +1 or -1. A symbol can have a disparity of +2 (six ones and four zeros), 0 (five ones and five zeros) or -2 (four ones and six zeros). If the disparity of a new symbol is anything other than +2, 0 or -2 it is invalid. When a new symbol arrives its disparity is calculated based on the current running disparity plus the disparity of the new symbol. The possible results are (running disparity + new symbol disparity):

(+1) + (+2) = +3 which is invalid

(+1) + (0) = +1

(+1) + (-2) = -1

(-1) + (+2) = +1

(-1) + (0) = -1

(-1) + (-2) = -3 which is invalid

When an invalid disparity arises it is an indication that something has gone wrong with the link and the link needs to be re-initialised. The running disparity can be tracked as soon as link is initialised.

When the 8B/10B encoder/decoder is separated into 5B/6B and 3B/4B encoders/decoders the above rules apply to both encoders/decoders. The disparity of each sub-code must be +2, 0 or -2 and the running disparity at the end of each encoder/decoder must be +1 or -1.

Table A-4 and Table A-5 show how errors can be captured by monitoring for invalid codes and disparity errors. The transmitter 5B/6B and 3B/4B codes are shown followed by the running disparity (+1 or -1) after the code has been sent. The initial running disparity is -1 in both examples. In Table A-4 a single bit error converts the D00.0 character sent into a code whose 4B component does not appear in the 3B/4B coding table and has a disparity of -4. This is immediately detected as a coding error.

Table A-4 Detection of error by invalid code

Character	Transmitted				Received			
D00.0	100111	+1	0100	-1	100111	+1	0000	-3 ERROR

In Table A-5 an error occurs in the first line with the D08.1 character being changed to the D05.1 character. D05.1 is a valid character so goes undetected. The running disparity should however be positive but because of the error it is negative. The characters that follow have neutral disparity so the running disparity remains unchanged and no error is detected. Eventually a character, D15.1, is sent which does not have neutral disparity. At the transmitter the running disparity is negative prior to D15.1 so that character is encoded as 101000 1001 which has negative disparity. When this is received at the receiver the negative disparity causes an error because the running disparity there is already negative. The error has been caught by disparity but several characters were sent before the error became apparent.

The receiver should look out for both invalid characters and disparity errors. It is also important that a CRC code is added to each packet sent to ensure that any error in a packet is detected.

Table A-5 Detection of error by invalid disparity

Char	Transmitted				Received				Char
D08.1	111001	+1	1001	+1	101001	-1	1001	-1	D05.1
D09.1	100101	+1	1001	+1	100101	-1	1001	-1	D09.1
D10.1	010101	+1	1001	+1	010101	-1	1001	-1	D10.1
D11.1	110100	+1	1001	+1	110100	-1	1001	-1	D11.1
D12.1	001101	+1	1001	+1	001101	-1	1001	-1	D12.1
D13.1	101100	+1	1001	+1	101100	-1	1001	-1	D13.1
D14.1	011100	+1	1001	+1	011100	-1	1001	-1	D14.1
D15.1	101000	-1	1001	-1	101000	-3 ERROR			D15.1

A.3 Serialisation and De-Serialisation

Serialisation is the conversion of a parallel data stream into a serial one. The parallel 10-bit data word is loaded into a shift register and then shifted out using a transmit clock signal to drive the shift register. A new character has to be loaded into the parallel input of the shift register as soon as the previous 10-bit character has been shifted out to prevent a gap in the serial data.

De-serialisation is the opposite of serialisation. The serial data is shifted into a shift register using a receive clock (also called a bit clock). Recovery of the receive clock from the transmitted serial data stream is described in section A.4 Once a full 10-bit character has been shifted into the shift register it is read out in parallel. The 10-bit character must be read out at the correct point in the serial data stream *i.e.* when a complete new 10-bit character is in the shift register. Character synchronisation is described in section A.5.

A.4 Receive Clock Recovery

Recovery of the receive clock (bit clock) from the received serial data stream is done using a phase-locked loop (PLL). A typical phase-locked loop is shown in Figure A-4.

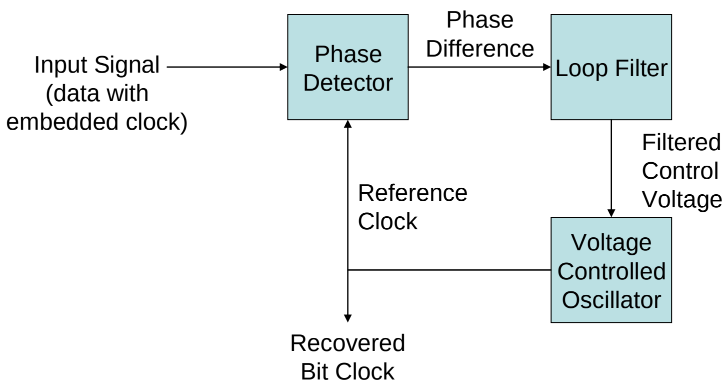

Figure A-4 Typical Phase-Locked Loop

The phase of the incoming data stream is compared to the phase of a reference clock signal. The detected phase difference is filtered, removing noise and providing an average phase difference. The filtered phase difference is used to control the frequency of the reference clock. If there is a positive phase difference with edges in the data stream occurring before edges in the reference clock then the reference clock frequency must be increased so that it catches up with the data stream edges. If there is a negative phase difference then the reference clock is occurring too early so must be slowed down, in which case the reference clock frequency is reduced.

When locked so that there is no phase difference, the reference clock can be used to recover the data-bits from the serial stream.

High frequency phase locked loops are normally implemented using a voltage controlled oscillator and an analogue loop filter.

The time taken for a PLL to lock onto a signal is dependent upon the design of the PLL and the difference between the input bit stream phase and the PLL reference clock phase. Typically it takes at least 5000 edges in the bit stream for a PLL to lock although it can be substantially longer for some PLL designs.

A.5 Symbol Synchronisation

Symbol synchronisation is necessary in the receiver to separate out each symbol from the received bit stream. To do this is it necessary to identify where a symbol starts, after that each individual symbol can be separated by simply counting 10-bits for each symbol. Identifying the start of a symbol is used using the 8B/10B Comma bit sequences. Comma sequences are unique seven bit sequences:

- Plus Comma 0011111
- Negative Comma 1100000

An illustration of symbol synchronisation using a plus comma is shown in Figure A-5. The start of the next character occurs on the fourth bit after the end of the detected plus comma.

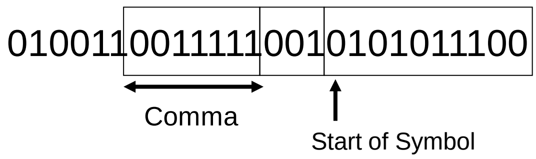

Figure A-5 Symbol Synchronisation Using a Plus Comma

Note that SpaceFibre uses all 10-bits of the commas for symbol synchronisation, rather than just 7-bits.

There are two principal means of performing symbol synchronisation. The first method, shown in Figure A-6, performs the symbol synchronisation after de-serialisation, while the second method, illustrated in Figure A-7, does it during de-serialisation.

Figure A-6 shows the received bit stream being fed into the de-serialising shift register. As soon as ten bits have been received the de-serialised data is loaded into a register. The exact position of the ten bits in the data stream is not important. After a further ten bits have been received the data in the register is loaded into a second register and the de-serialised data is loaded into the first register. The 20 bits in these two registers are examined for a possible comma, using the comma detect circuitry. The combinatorial logic in the comma detect circuitry outputs a Comma Detect signal when a comma is found and registers the position of the start of the next character. The Start of Symbol is used to drive a 20:10 multiplexer which select the ten-bits of a character from the 20-bits in the two registers.

With this approach comma detection is done at a clock rate of one tenth of that of the bit stream. The comma detect circuitry has to simultaneously look for a comma in ten possible bit positions, requiring 10 correlators.

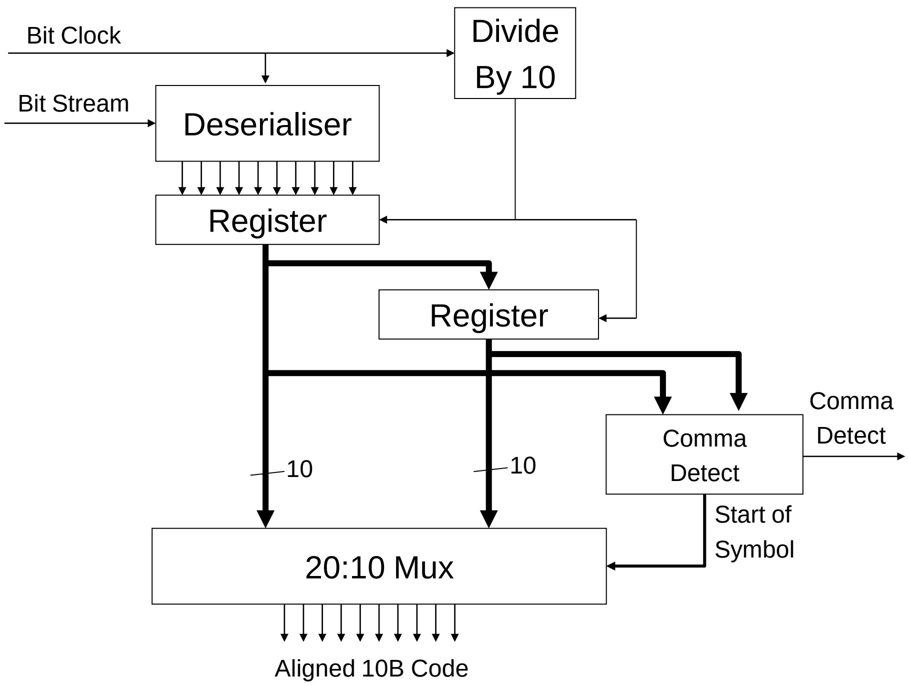

Figure A-6 Symbol Alignment After De-serialisation

The other approach, shown in Figure A-7, performs comma detection at the bit clock rate. The bit stream is fed into a 10-bit shift register (de-serialiser). The 10-bit parallel output from the shift register fed to a 10-bit character register and to a Comma Detect circuit. The Comma Detect circuit looks for a Comma in the last seven bits of the shift register (*i.e.* the first seven bits to enter the shift register). When a comma is detected the data in the shift register are loaded into the data register. A bit counter is used to count the 10-bits in each character, loading the data register from the shift register every 10 bits. The comma detect circuit resets the counter, re-synchronising the bit counter and forcing the data in the shift register to be loaded into the data register.

This approach requires the comma detect circuitry to operate at the rate of the bit clock and needs a high-speed bit counter. The amount of circuitry is significantly less that the other approach.

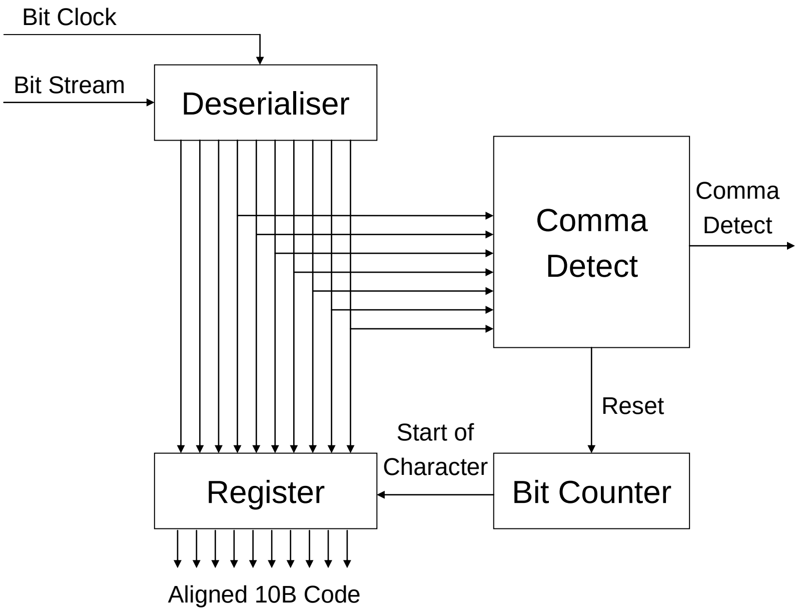

Figure A-7 Character Alignment During De-serialisation

A.6 Receive Elastic Buffer

The two ends of a link are both expected to operate at the same frequency. In practice, however, there will be slight differences in the clocks at the two ends of the link. This can cause receive buffer overflow or under-run problems unless the difference in the two clock speeds is compensated for. This is achieved using a Receive Elastic Buffer and associated SKIP characters.

The receive clock (bit clock) is recovered from the incoming bit stream, so is at the same frequency as the transmit clock at the other end of the link. After de-serialisation and character synchronisation the incoming data must be transferred from the receive clock domain to the local system clock domain. In passing between these two clock domains, slight differences in the clock frequencies must be accommodated. This is achieved using the Receive Elastic Buffer.

The normal situation with a Receive Elastic Buffer is illustrated in Figure A-8. Data is written into the buffer using a write pointer which operates at the receive character rate (RXRECCLK). It is read out by read pointer which operates at the user system character rate (RXUSRCLK).

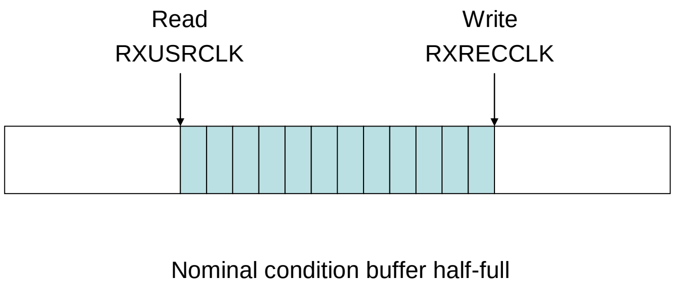

Figure A-8 Receive Elastic Buffer - Nominal Condition

Any difference between the two clock frequencies is compensated for using a special character, SKIP, inserted every so often into the data stream

If the RXUSRCLK is faster than RXRECCLK the buffer will slowly empty. When the buffer is less than half full, implying that RXUSRCLK is faster than RXRECCLK, extra SKIP characters are added to the Receive Elastic Buffer. This may be done when a SKIP character is read out of the buffer, by simply not incrementing the read pointer, so that the SKIP character will be read a second time. The effect is to add an extra SKIP character to the data stream, temporarily slowing down the RXUSRCLK to compensate for it being faster than RXRECCLK. This is illustrated in Figure A-9.

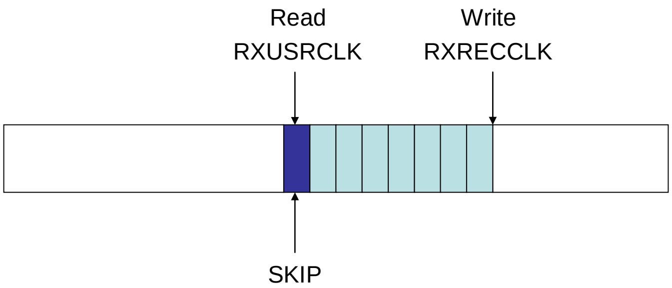

Figure A-9 Receive Elastic Buffer Emptying

If RXUSRCLK is slower than RXRECCLK then the Receive Elastic Buffer will slowly fill up. When the buffer is more than half full, SKIP characters are skipped. This is done by incrementing the read pointer past a SKIP character, *i.e.* if after reading a character, the next character to be read is a SKIP character, it is ignored and the read pointer is moved to point to the following character instead. The effect is to remove SKIP characters from the buffer, temporarily speeding up the RXUSRCLK to make up for the fact that it is slower than RXRECCLK. This is shown in Figure A-10. Note that the SKIP operation requires the elastic buffer to know in advance that a SKIP is present in the buffer without reading it otherwise the SKIP operation will have no effect.

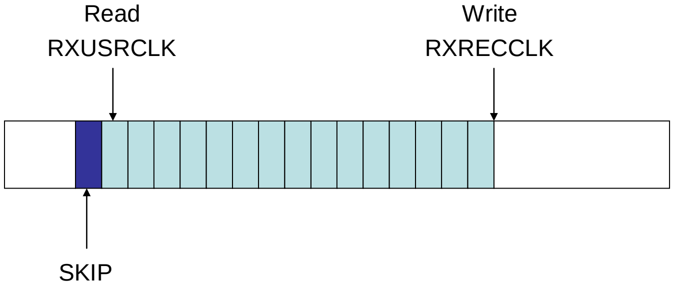

Figure A-10 Receive Elastic Buffer Filling Up

For the Receive Elastic Buffer to work properly there must be sufficient SKIPs in the data stream, so that they can be remove if necessary. The frequency of SKIPs depends on the size of the elastic buffer and the maximum frequency difference between RXUSRCLK and RXRECCLK.

Assume that the nominal operational frequency is F symbols per second and that the maximum clock difference, is D Hz, then the time, T, taken for the elastic buffer to have one symbol too many or one symbol too few is given by:

T = 1 / (Receive Clock Frequency – User Clock Frequency)

T = 1/ ((F+D) – (F-D))= 1/ (2D)

In this time the number of symbols sent is

N = (F+D).T

which is approximately

N ≈ F/(2D)

since F is much greater than D.

Now D/F is the maximum clock drift, so

N ≈ 1/(2P)

Where P is the maximum drift in the clock.

For a ±100 ppm maximum clock drift, which is readily achievable using crystal oscillators, D, is 10^{-4}, and the number of symbols sent before the elastic buffer is one symbol out is 5000. A SKIP symbol must thus be sent every 5000 symbols to prevent the Elastic buffer from ever being more than one symbol out. This is the case independent of the size of the symbols.

In SpaceFibre the receive elastic buffer stores control and data words rather than individual symbols. The SKIP control word is therefore four symbols lone, i.e. one word long.

Annex B (informative) Example of SpaceFibre CRC implementation

B.1 Overview

In this example implementations of the CRC used by SpaceFibre are provide in VHDL and C-code.

B.2 VHDL implementation of SpaceFibre CRC

Bibliography

ECSS-ST-S-00	ECSS system - Description, implementation and general requirements
ECSS-E-ST-50	Space engineering - Communications
ECSS-E-HB-50	Space engineering - Communications guidelines
ECSS-E-ST-40	Space engineering - Software
http://www.spacewire.esa.int	SpaceWire website