ATM Forum
Technical Committee

UTOPIA Level 4

AF-PHY-0144.001

March, 2000

The ATM Forum
Worldwide Headquarters
2570 West El Camino Real, Suite 304
Mountain View, CA 94040-1313
Tel: +1-650-949-6700
Fax:+1-650-949-6705

Contact:

John R. Mick, Chair
Phone: +1 408 654 6721
Email: mick@idt.com

Dave King, Editor
Phone: +1 781 455 5832
Email: dave.king@gd-cs.com

Table of Contents

List of Figures

List of Tables

Blank Page

1 INTRODUCTION

1.1 Purpose

This specification describes a new generation of the Universal Test & Operations PHY Interface for ATM (UTOPIA) data path interfaces. The Utopia Level 4 specification provides similar functionality as the Utopia Level 1, 2 and 3 interfaces, but will cater for higher data rates. Familiarity with the Level 1, 2 and Level 3 Utopia specifications would be beneficial to the reader but is not required.

While the Utopia Level 4 specification will build on concepts described in Utopia Level 1, 2 and 3 specifications, a large number of changes and improvements have been made and this specification can be used on its own.

1.2 Scope

The Utopia Level 1, 2 and 3 specifications are based on common bus transfer techniques that allow data transfer rates up to 3.2Gbps. These interfaces are based on the host providing timing and control. The need to pass 10Gbps over the Utopia Level 4 however, dictates the need for new control techniques and a new IO specification. This specification is for an interface that differs from its predecessors in the following ways:

1) 32, 16 or 8 bit Low Voltage Differential Signaling (LVDS) bus.
2) Symmetric transmit/receive bus structure.
3) Nominal transfer rates of 10Gbps.
4) Inband control of cell delimiters and flow control.
5) Use of source synchronous clocking.
6) Support of variable length payloads or packet systems.

1.3 Abbreviations and Acronyms

ADR	Address
ATM	Asynchronous Transfer Mode
CLP	Cell Loss Priority
EOP	End Of Packet
EXT	Extended
FIFO	First In First Out
FLC	Flow Control
GFC	Generic Flow Control
HEC	Header Error Control
I/F	Interface

LVDS	Low Voltage Differential Signaling
PHY	OSI Physical Layer
PTI	Payload Type Identifier
SDH	Synchronous Digital Hierarchy
SOCP	Start of Cell or Packet
SONET	Synchronous Optical Network
UDF	User Defined Fields
UTOPIA	Universal Test & Operations PHY Interface for ATM
VCI	Virtual Channel Identifier
VPI	Virtual Path Identifier

Real differences between SONET and SDH (other than terminology) are minor; for the purposes of this specification, they are inconsequential or irrelevant.

For the convenience of the reader, equivalent terms are listed below:

SONET	**SDH**
SPE	VC
STS-SPE	Higher Order VC (VC-3/4/4-Nc)
STS-1-SPE	VC-3
STS-3c frame	STM-1 frame, AU-4
STS-3c-SPE	VC-4
STS-3c payload	C-4
STS-12c/48c/192c frame	STM-4/16/64 frame, AU-4-4c/16c/64c
STS-12c/48c/192c-SPE	VC-4-4c/16c/64c
STS-12c/48c/192c payload	C-4-4c/16c/64c

1.4 Configurations

As in the Utopia Level 3 Specification, this interface supports point-to-point, high-speed interconnections. Examples are shown in Figure 1.

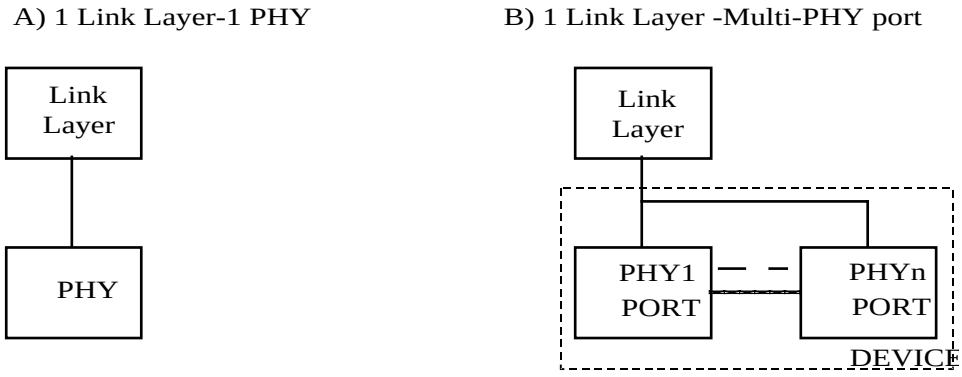

Figure 1: Link Layer to PHY Interconnect

2 FUNCTIONAL DESCRIPTION

The Asynchronous Transfer Mode (ATM) or packet system physical interface, as specified by this proposal, has the following characteristics and features:

1) Supports the equivalent capacity of a Synchronous Optical Network (SONET) OC-192 (9.95328 Gb/s) signal with single or multiple OSI Physical Layer (PHY) ports on one PHY device.
2) Provides addressing support for payloads channeled down to SONET STS-1 with addressing support for even deeper channeling.
3) Minimizes the number of signals required on the interface by moving all control functions in-band.
4) Provides a future upgrade path for rates higher than OC-192 and for additional control functions, including user-defined functions.
5) Supports interconnections across motherboard, daughterboard and backplane interfaces.
6) Completely symmetrical interface in Tx and Rx directions for easier design and testing (e.g., simple loopback for fault isolation), and wider applicability.
7) Simple and efficient flow-control allows handshaking with up to 24 concurrent channels while providing flow control information in one clock cycle (no polling required).
8) Uses source synchronous clocking methodology for easier board and system design.

The bus width can be 32, 16 or 8-bits. At least one of these widths must be supported. The basic interface operates at up to 415MHz[1]. In addition to data, the only other signals are "clock" and "ctrl", which indicate whether a data or control word is being transferred. Flow control, addressing and other control functions are all performed by control transfers over the data bus.

While a common application for such a Packet System interface is envisaged to be data transfer between a PHY entity and a link layer entity, the protocol is not limited to this configuration. Owing to its symmetry, the protocol is also particularly suitable for link layer transport across peer entities. Where transport is between a PHY entity and a link-layer entity, the Transmit (Tx) direction shall refer to the Link Layer-to-PHY direction of transmission and the Receive (Rx) direction shall refer to the PHY-to-Link layer direction of transmission.

[1] 415 MHz is just over 8 times the OC-1 clock rate, which can be a convenient system frequency in SONET applications.

Blank Page

3 SIGNALS

Table 1 and Table 2 describe the signals and functions necessary to support the interface.

Table 1: ATM/Link to PHY Layer Signals

Name	No. Pins	Function
tx_clk	2	Source clock from ATM to PHY
tx_ctrl	2	Indicates if tx_data is control or data information
tx_data [N-1:0]	2*N	N bit data bus (where N=32 or 16 or 8)

Table 2: PHY to ATM/Link Layer Signals

Name	No. Pins	Function
rx_clk	2	Source clock from PHY to ATM
rx_ctrl	2	Indicates if rx_data is control or data information
rx_data [N-1:0]	2*N	N bit data bus (where N=32 or 16 or 8)

The 32-bit interface requires 68 pins in each direction for a total of 136 pins full duplex. The 16-bit interface requires 36 pins in each direction for a total of 72 pins full duplex. The 8-bit interface requires 20 pins in each direction for a total of 40 pins full duplex. The symmetry of UTOPIA Level-4 is illustrated in Figure 2.

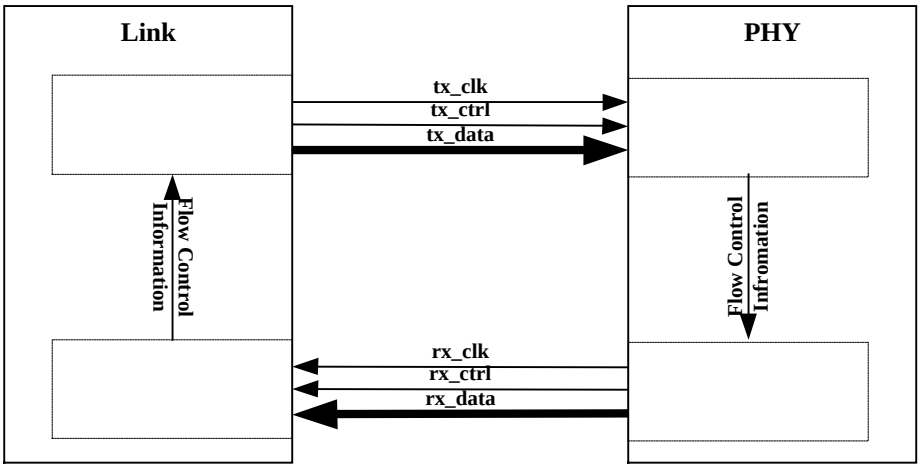

Figure 2: Signal Definition

Blank Page

4 DATA FORMAT

4.1 General Format Description

The interface control bits rx_ctrl and tx_ctrl are used to differentiate the associated 32-bits of data as being either control information or cell/packet data. The definition of these bits is shown in Table 3.

Table 3: Tx_ctrl and Rx_ctrl Values

Transfer Ctrl	Rx/tx_ctrl	Description
DATA	0	ATM header/Packet or cell data
CTRL	1	Control information

The ATM payload format is shown in Table 4. The data consists of 52 bytes because the Header Error Control (HEC) and User Defined Fields (UDF) fields are not represented as data. A complete payload transfer over a 32-bit bus takes 13 clock cycles. Bus widths of 16 and 8 are also supported (see section 4.3 and 4.4) but all data transfers are a logical 32-bit width.

Table 4: ATM Cell Data Format (rx/tx_ctrl=0)

RX/TX_CTRL	Rx_data/tx_data	Description
DATA	Header_1 to Header_4	1st transfer of ATM cell data
DATA	Payload_1 to Payload_4	2nd transfer of ATM cell data
DATA	Payload_5 to Payload_8	3rd transfer of ATM cell data
*	*	*
DATA	Payload_41 to Payload_44	12th transfer of ATM cell data
DATA	Payload_45 to Payload_48	13th transfer of ATM cell data

The CTRL transfer is used to transmit multi-PHY address information, flow control information and start-of-cell information. Table 5 shows the possible values of **rx/tx_data** when **rx/tx_ctrl** = ‘1’.

Table 5: Control Format (rx/tx_ctrl=1)

Name	Rx/Tx_data Bit Definitions									Function
	31:30	**29**	**28**	**27:26**	**25**	**24**	**23:8**		**7:0**	
ADR	00	SOCP	EOP	SIZE	P	0	CONG [7:0]	REL [7:0]	ADR [7:0]	Next PHY address, Incremental Congestion, SOCP, EOP
		SOCP	EOP	SIZE	P	1	*		ADR [7:0]	Next PHY address, SOCP, EOP
FLC	01	00		BLK [3:0]			FULL [23:0]			Flow control, Idle
		01-11		*						Reserved for future definition
EXT	10	0000			*			PRTY_C	PRTY [7:0]	Parity
		0001			*				ABORT [7:0]	Packet Abort
		0010			*				HEC [7:0]	ATM Cell HEC field
		0011-1111			*					Reserved for future definition
UDF	11	**								Reserved for user defined fields

* Reserved for future use and shall be set to zero when transmitted, ignored when received.
** Reserved for users.

ADR: Each data transfer (both cell and packet data) must be preceded by an ADR word. The ADR word is used to indicate the start of the next cell/packet and the end of the preceding packet. ATM Cell data is not required to be followed by an ADR word. The format of the ADR word as shown in Table 5 indicates support for up to 255 PHY channels. The ADR word also has the following properties:

a) Start of Cell or Packet (SOCP*)* indicates the start of a cell or packet and End of Packet (EOP) indicates the end of a packet. Start of Packet (SOCP) (with a new address) may be indicated concurrently with EOP.
b) The SIZE field indicates valid bytes in data transfers (one 32 bit logical word) of between 1 and 4 bytes. This field is valid only with EOP='1' indicating that the previous logical word (32 bits) was the final data transfer of a packet. The number of valid bytes 1 (only the MSB is valid), 2, 3 or 4 are indicated by the SIZE field values of '01', '10', '11' or '00' respectively.
c) The P bit indicates whether the data being transferred is an ATM cell (logic '0') or a packet (logic '1').
d) The two incremental congestion fields, CONG[7:0] and REL[7:0], are used for flow control by indicating changes in full conditions of one of the 255 possible ports. The congestion field, CONG[7:0], indicates that the port addressed at CONG[7:0] has just undergone a full condition. The release field, REL[7:0], indicates that port addressed at REL[7:0] has just come out of a full condition. If no congestion change has occurred, a CONG[7:0] or REL[7:0] value of all '1's is sent. Bit 24 of the ADR word may disable the incremental congestion field in cases where backpressure is not needed.
e) The ADR[7:0] contains the address for the next data transfer. If no data is available and an idle Flow Control (FLC) is to be transferred next, the ADR[7:0] field shall be set to all '1's. An ADR

word prior to idle FLC transfers allows end of cell and EOP closure to the last valid cell or packet transfer.

When a channel is provisioned for ATM operation, the EOP and SIZE fields are not used and shall be set to zero: the completion of the 52nd payload byte transfer automatically marks the end of the cell.

FLC: When received, Flow Control words tell the local transmitter (see Figure 2) to stop sending DATA within 40 word transfers (as defined in sections 4.2, 4.3, and 4.4). They may be interspersed among DATA and other CTRL transfers. Section 6 defines the port mapping for this field.

a) Indications may simultaneously be provided for 24 channels on FULL[23:0].
b) The block bits, BLK[3:0], select which group of 24 channels the FULL[23:0] signals are reporting, supporting the 255 channels as described in section 6.2. The use of the reserved bits permits channeling to be increased further.

The FLC transfers shall also be used (with none of the full bits asserted if none apply) for idle transfers when no other data or control information needs to be transferred to give faster update of flow control status.

EXT: Used for extended functions for future standardization. The presently defined uses are for the Parity and Abort Packet words.

a) Parity is calculated over all the bits transferred since the last parity. The parity field, PRTY[7:0], would be an 8-bit-interleaved odd parity over all the bytes in the 32-bit transfers. Note that the calculation of parity is identical for the 32, 16 and 8-bit modes, (i.e., it is still based on 32-bit logical words). Parity over the rx/tx_ctrl signal, PRTY_C, is also generated. Figure 3 illustrates how the parity shall be calculated beginning at the last parity control word.
b) Packet Abort is sent to indicate that the transmitter has determined that the last packet from address ABORT[7:0] should be dropped.
c) The ATM HEC is not required to be transferred as part of Utopia Level 4. However, users who wish to implement it should use this EXT field. The nominal location for this extension would be immediately following the Virtual Path Identifier (VPI)/Virtual Channel Identifier (VCI) data transfer.

ctrl	MSb															LSb														LSb
	MSByte																								LSByte							
CD1=1	A1	B1	C1	D1	E1	F1	G1	H1	A2	B2	C2	D2	E2	F2	G2	H2	A3	B3	C3	D3	E3	F3	G3	H3	A4	B4	C4	D4	E4	F4	G4	H4
CD2	A5	B5	C5	D5	E5	F5	G5	H5	A6	B6	C6	D6	E6	F6	G6	H6	A7	B7	C7	D7	E7	F7	G7	H7	A8	B8	C8	D8	E8	F8	G8	H8
CD3	A9	B9	C9	D9	E9	F9	G9	H9	A10	B10	C10	D10	E10	F10	G10	H10	A11	B11	C11	D11	E11	F11	G11	H11	A12	B12	C12	D12	E12	F12	G12	H12
CD4	A13	B13	C13	D13	E13	F13	G13	H13	A14	B14	C14	D14	E14	F14	G14	H14	A15	B15	C15	D15	E15	F15	G15	H15	A16	B16	C16	D16	E16	F16	G16	H16
				:									:							:									:		
CDN	AN-3	BN-3	CN-3	DN-3	EN-3	FN-3	GN-3	HN-3	AN-2	BN-2	CN-2	DN-2	EN-2	FN-2	GN-2	HN-2	AN-1	BN-1	CN-1	DN-1	EN-1	FN-1	GN-1	HN-1	AN	BN	CN	DN	EN	FN	GN	HN

Parity Word:			< =PRTY_C= >	< ==== PRTY[7:0] ==== >							
31:30	29:26	25:9	8	7	6	5	4	3	2	1	0
CDN+1 =1	10	0000	Reserved	=XNOR[CD1:CDN]	=XNOR[A1:AN]	=XNOR[B1:BN]	=XNOR[C1:CN]	=XNOR[D1:DN]	=XNOR[E1:EN]	=XNOR[F1:FN]	=XNOR[G1:GN]	=XNOR[H1:HN]

Figure 3: Calculation of Parity

UDF: Reserved for User Defined Fields. The frequency and location of UDF transfers relative to other transfers are not defined by this specification, and may be user-defined.

4.2 32-Bit, 10Gbps Interface Mode (Optional)

All DATA and CTRL word formats are as described in section 4.1, with a logical word size of 32-bits. Hence, a logical word transfer takes 1 clock cycle as shown in Figure 4. Any instance of the interface operates independently, with its own address space, flow control, clock and control pins. At the nominal clock rate, this interface can transfer up to the full 10Gbps.

	Most Significant Byte [31:24]			Least Significant Byte [7:0]
Data1	A	B	C	D
Control	E	F	G	H
Data 2	I	J	K	L
Data 3	M	N	O	P

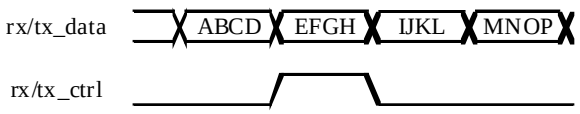

Figure 4: Transfer of 32-Bit Mode

The **rx/tx_ctrl** signal is set to ‘1’ for the first clock cycle of the 32-bit logical control word, and is ‘0’ otherwise. This permits unambiguous 32-bit word delineation on the 32-bit bus as shown in Figure 4. If the data were an ATM cell, byte A would represent the Generic Flow Control (GFC) and the most significant bits of the VPI. Byte D would represent the least significant bits of the VCI, the Payload Type Identifier (PTI) and Cell Loss Priority (CLP) bits.

4.3 16-Bit, 5Gbps Interface Mode (Optional)

All DATA and CTRL word formats are as described in section 4.1, with a logical word size of 32-bits. Hence, a logical word transfer takes 2 clock cycles as shown in Figure 5. Any instance of the interface operates independently, with its own address space, flow control, clock and control pins. At the nominal clock rate, this interface can transfer up to 5Gbps.

	Most Significant Byte [31:24]			Least Significant Byte [7:0]
Data1	A	B	C	D
Control	E	F	G	H
Data 2	I	J	K	L
Data 3	M	N	O	P

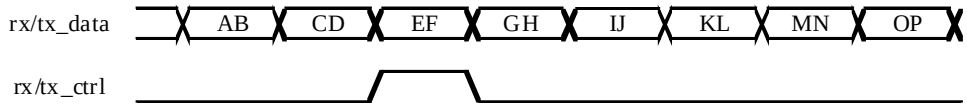

Figure 5: Transfer of 16-Bit Mode

The **rx/tx_ctrl** signal is set to '1' for the first clock cycle of the 32-bit logical control word, and is '0' otherwise. This permits unambiguous 32-bit word delineation on the 16-bit bus. Rx/Tx_data(31:16) is transmitted first, followed by rx/tx_data(15:0) shown in Figure 5. If the data were an ATM cell, byte A would represent the GFC and the most significant bits of the VPI. Byte D would represent the least significant bits of the VCI, PTI and CLP bits.

4.4 8-Bit, 2.5Gbps Interface Mode (Optional)

All DATA and CTRL word formats are as described in section 4.1, with a logical word size of 32-bits. Hence, a logical word transfer takes 4 clock cycles as shown in Figure 6. Any instance of the interface operates independently, with its own address space, flow control, clock and control pins. At the nominal clock rate, this interface can transfer up to 2.5Gbps.

	Most Significant Byte [31:24]			Least Significant Byte [7:0]
Data1	A	B	C	D
Control	E	F	G	H
Data 2	I	J	K	L
Data 3	M	N	O	P

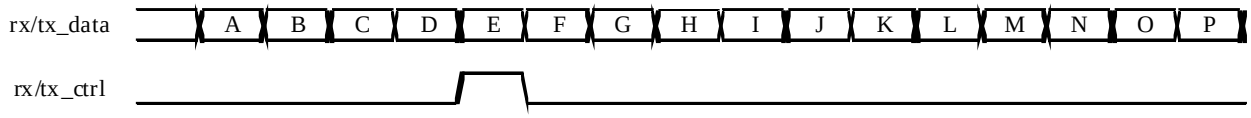

Figure 6: Transfer of 8-Bit Mode

The **rx/tx_ctrl** signal is set to '1' for the first clock cycle of the 32-bit logical control word, and is '0' otherwise. This permits unambiguous 32-bit word delineation on the 8-bit bus. Rx/Tx_data(31:24) is transmitted first, followed by rx/tx_data(23:16), rx/tx _data(15:8) and finally rx/tx _data(7:0) shown in Figure 6. If the data were an ATM cell, byte A would represent the GFC and the most significant bits of the VPI. Byte D would represent the least significant bits of the VCI, PTI and CLP bits.

Blank Page

5 DATA TRANSFER

A SONET OC-192 interface provides payload at a peak rate of 311.04 MHz on a 32-bit bus; when the SONET overhead is subtracted, the average rate is 299.52 MHz. This allows approximately 25% of the bandwidth of a 400 MHz interface available for CTRL data. If the interface is transferring ATM, the 52-byte payload takes 13 clock cycles, allowing an additional 4.33 clock cycles during each cell transferred for control. A single-cycle ADR transfer must be done before each ATM cell to indicate the start of a cell and its PHY address. When packets are transferred, each packet needs, as a minimum, one CTRL transfer for address, SOCP and EOP information. This maximizes the bandwidth that is available for other control functions such as flow control, parity, etc.

5.1 Protocol Rules

5.1.1 ATM Rules

The following rules define the protocol for ATM cell transfers:

1. All ATM cells are 52 bytes of data transfers.
2. An ADR transfer with SOCP set must occur before data transfers for that address indicating the start-of-cell.
3. The 52nd byte transferred indicates the end-of-cell for the cell currently in progress.
4. Data within a cell transfer shall not be interrupted except by EXT, FLC and UDF transfers.
5. There are two cells of receive buffering and FLC transfers stop data only between cells.
6. FLC words, and not Incremental Congestion Indications, shall be used to indicate congestion changes when address changes are not being done.

5.1.2 Packet Rules

The following rules define the protocol for packet transfers:

1. An ADR transfer with SOCP=’1’ must occur before data transfers for that address indicating the start-of-packet.
2. In packet mode, an ADR transfer with EOP must occur to indicate the end-of-packet for the previously active address. It may occur concurrently with an ADR/SOCP to indicate the start of the next packet. If there is no other packet awaiting transfer, the value of the address field is set to all ‘1’s, since it will be superseded by the actual address of the following transfer. The control word containing EOP also indicates the number of significant bytes in the final data word transferred for that address.
3. Packets shall be sent in blocks of no less than 64 contiguous data bytes. A block size of 64 bytes shall always be supported. Larger block sizes may also be supported. There are two exceptions to

the minimum block size transfer rule. First, if less than 64-bytes remain in the packet, only the remaining words are transferred. The second exception occurs if the packet is to be aborted. .

4. Data within a packet block transfer shall not be interrupted except by EXT, FLC and UDF transfers.
5. In packet mode, addresses may be changed mid-way through a packet in compliance with the minimum block transfer; if this change is not accompanied by an EOP, it suspends data transfer for the previously active address and resumes it the next time that address is made active.
6. A second packet with the same address shall not interrupt a packet transfer. The Abort Packet extension shall be used to delete a packet.

5.1.3 Protocol for Receive State Machine

The protocol state machine for each channel's receiver has four states as follows:

A. **IDLE** - no cell/packet is currently being transferred to this channel.

B. **START** - the next data transfer will be the beginning of a cell/packet for this channel.

C. **ACTIVE** - data transfer of a cell/packet for this channel has begun.

D. **PAUSE** - data transfer has been switched to another channel in mid transfer (packet only).

Figure 7 shows the state transition diagram for each channel.

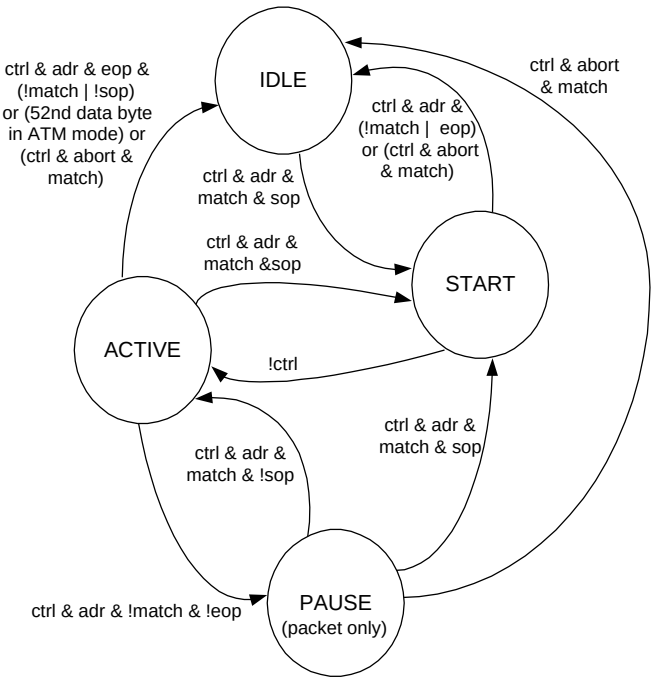

Figure 7: Per Channel Receiver State Diagram

Each channel is looking at ADR transfers for addresses that match its own. In the state diagram, ctrl means that the transfer type is “control,” adr means that the control type is “address.” The term match will be used for ADR or Abort EXT transfers that match the channel, and !match will be used for those that don’t. The sop (!sop) and eop (!eop) indicate whether the SOCP and EOP bits are set (cleared), respectively.

Blank Page

6 ADDRESSING

6.1 Relationship between Port Numbers on PHY I/F and Addresses on Utopia 4 I/F

For transport of cells/packets between PHY layer and link layer entities, the port addresses shall, wherever possible, be fixed and determined by the natural port addressing of the PHY layer (this should eliminate the need for address map provisioning at the PHY end of the Utopia 4 interface). Where the nature of the interface does not permit determination of a natural port numbering order, the port number to address mapping shall be made programmable.

6.2 Mapping of Utopia 4 Addresses to Flow Control Encoding

Utopia 4 addresses shall be mapped to the BLK[3:0] and FULL[23:0] fields in the FLC control word shown in Table 5 as follows:

1. BLK[3:0] = int(address/24).
2. FULL[23:0] can encode 24 separate addresses within the same block. For an address in a block denoted by BLK[3:0]:
3. FULL[x] is set if (address % 24) = x (where % is the modulo operator, i.e., the remainder when address is divided by 24)

In this way, up to 24 ports can indicate their congestion status with a single FLC transfer.

Blank Page

7 FLOW CONTROL

Operation of flow control shall observe the following rules:

1. Each side of a Utopia 4 interface shall maintain a state for each active address that indicates if the buffer for that address is approaching its capacity for storing additional cells. Although this rule is stated in terms of a fixed buffer per address, it is not intended to preclude implementations with shared or variable buffers.
2. At periodic intervals, an FLC control word shall be sent to the far end indicating which address on the near end is approaching its capacity.
3. FLC words may be sent at any time, including between data transfers.
4. For up to 24 addresses, a single FLC transaction handles all addresses in one transaction. If there are greater than 24 addresses, it is theoretically possible that multiple FLC transactions may be needed. The frequency of FLC transmission, buffer size threshold, and operating frequency of the Utopia 4 interface all determine the maximum number of addresses for which FLC can be exercised simultaneously with deterministic worst-case behavior.
5. Incremental Flow Indicators as defined in Section 4, can indicate that a channel has undergone a congestion change. Note that a single ADR word can indicate congestion and a release for two different port channels. If no congestion change has occurred, a CONG[7:0] or REL[7:0] value of all 1’s is sent.
6. Since the receiver can accept congestion information from either FLC words or the incremental indicators in ADR words, the most recently received shall be considered the valid state of the channel.
7. FLC words, and not Incremental Congestion Indications, shall be used to indicate congestion changes when address changes are not being done.
8. After the receipt of a congestion condition, the transmitter shall cease DATA transmission to the congested channel within 40 data words. Data transfers to the channel may resume after the receipt of a congestion release indication.
9. In order to maintain proper port synchronization, it is recommended that an indication of every channels congestion status shall be transmitted at least every 2E20 clock cycles. Failure to do so is not considered an error condition.

7.1 ERROR CONDITIONS AND ERROR HANDLING

Table 6 shows the error conditions on the Utopia 4 bus protocol, and consequent recovery actions.

Table 6: Error Conditions and Recovery Action

Condition	Action
Missing SOCP (address change without SOCP and no suspended transfer for that packet)	All subsequent data words for that address ignored until following SOCP
Invalid (non-existent, not provisioned) ADR[7:0]	All subsequent data ignored until next valid address. All other ADR control bits shall be considered valid
Missing EOP (new SOCP received for same address)	Previous partial packet discarded or aborted. New SOCP considered valid for subsequent data transfer
Unexpected (too early) SOCP (ATM mode)	Previous partial cell discarded. SOCP considered valid for following cell
Missing SOCP (ATM mode, >52 bytes of data transferred)	Previous cell (up to 52nd byte) considered valid. All subsequent data words ignored till following SOCP
Premature address change (ATM mode, before 52 bytes transferred for previous address)	Previous partial cell discarded. New address effective for all subsequent data transfers
Invalid (non-existent, not provisioned) CONG[7:0]	Congestion set indication for [7:0] shall be ignored. All other bits of the ADR word including addresses ADR[7:0], REL[7:0] shall be considered valid.
Invalid (non-existent, not provisioned) REL[7:0]	Congestion release indication for [7:0] shall be ignored. All other bits of the ADR word including addresses ADR[7:0], CONG[7:0] shall be considered valid.
EXT with Abort for a port that has already received a complete packet	Ignore EXT word.
ADR with SOP following an ADR with SOP (other port address).	Ignore first ADR; if packet data for the port address of the first ADR is to be sent later, it must start with a new ADR. Congestion information of the ignored ADR word is considered valid.

8 TIMING

Differing from previous Utopia specifications where the ATM host is the source of both transmit and receive timing, this specification utilizes a source synchronous timing method as shown in Figure 8. The advantage of using the source synchronous clocking mechanism is that far end data recovery is much more easily achieved. There is no need to account for chip I/O delays or use Phase Lock Loop techniques to ensure clock/data alignment.

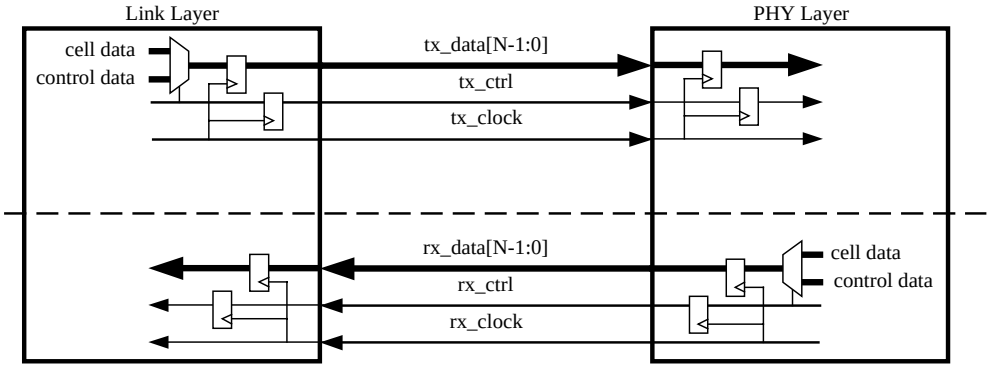

Figure 8: Source Synchronous Timing

The following figures are examples illustrating the cell and packet transfers that are compliant with this specification. As described above, the transmit signals and timing are identical to those of the receive.

Figure 9 illustrates a single cell transfer followed by a parity EXT control word and an idle/full control word. It should be noted that the full indication for the receive port 10 allows that 2 cells worth of First-In-First-Out (FIFO) depth is still available.

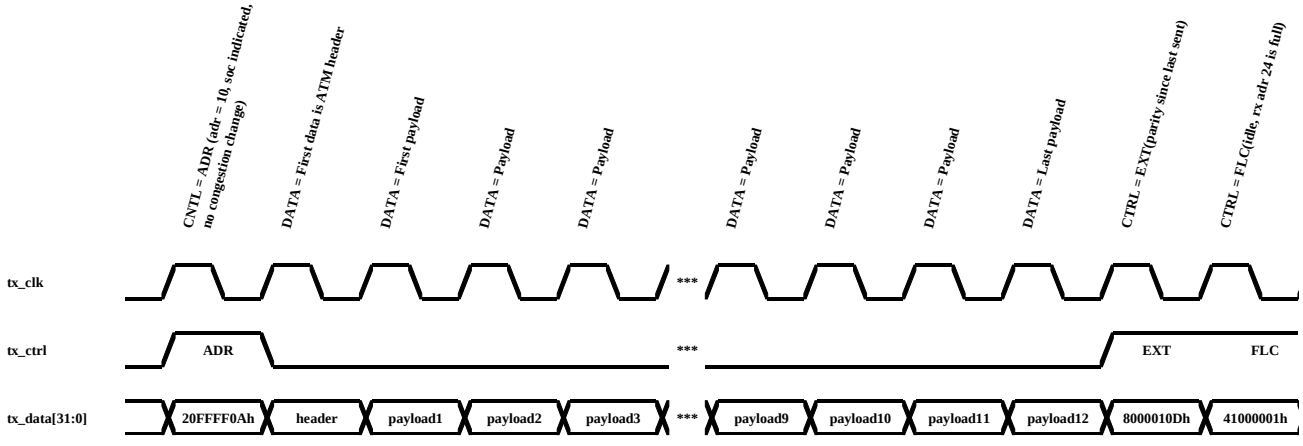

Figure 9: Single Cell Transfer to Port 10

Figure 10 illustrates back-to-back cell transfers to the same address while Figure 11 demonstrates back-to-back cell transfers to different addresses. These examples demonstrate the most efficient cell transfer of the interface.

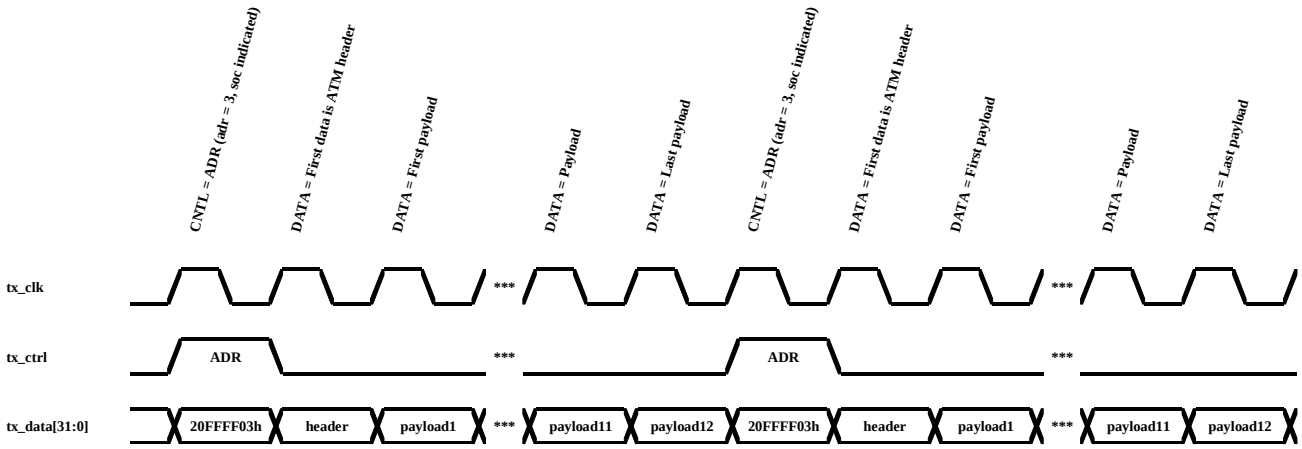

Figure 10: Back to Back Cell Transfer

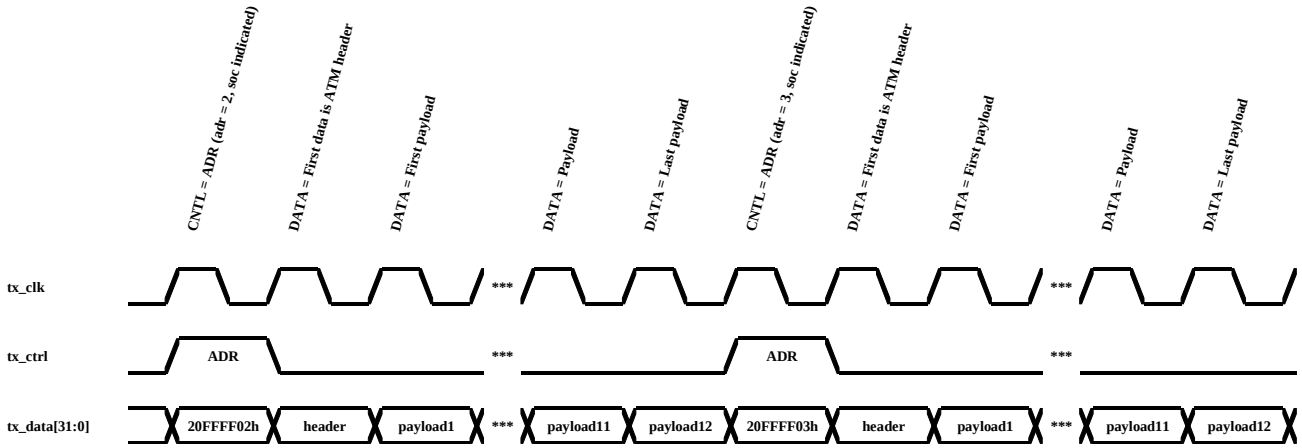

Figure 11: Back to Back Cell Transfer, Two Different Ports

Flow control can be intermixed while data is being transferred. Figure 12 illustrates this by showing how the receive bus's FIFO has indicated a full condition while transmitting a user cell.

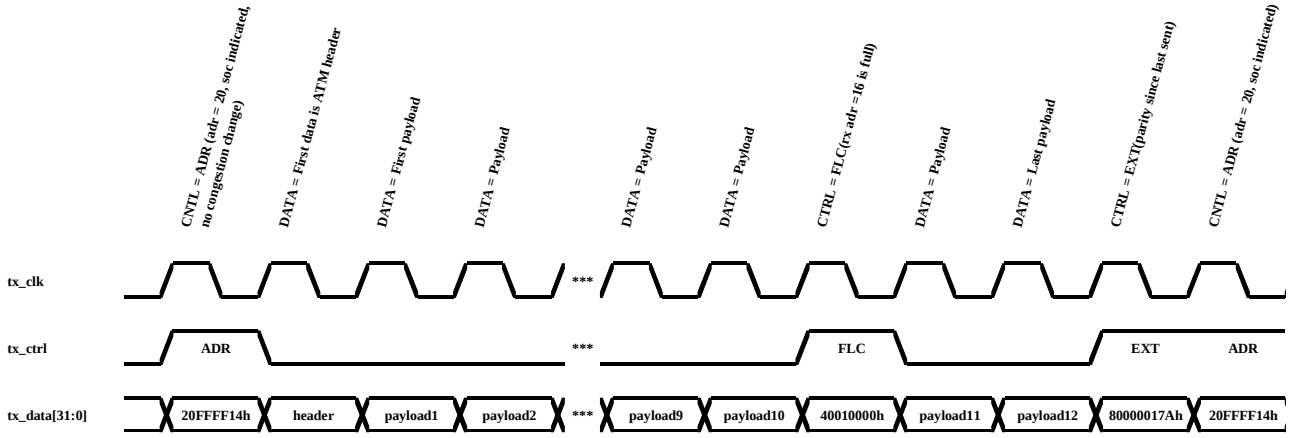

Figure 12: Cell Flow Control

Figure 13 illustrates how this interface could be used in an 8-bit configuration.

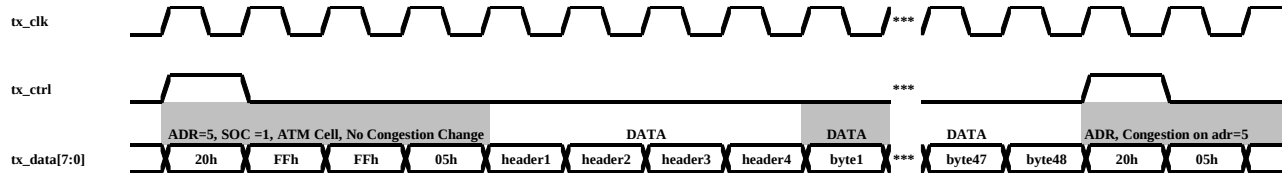

Figure 13: 8-Bit Mode Transfer

Figure 14 demonstrates the incremental congestion and idle transfers.

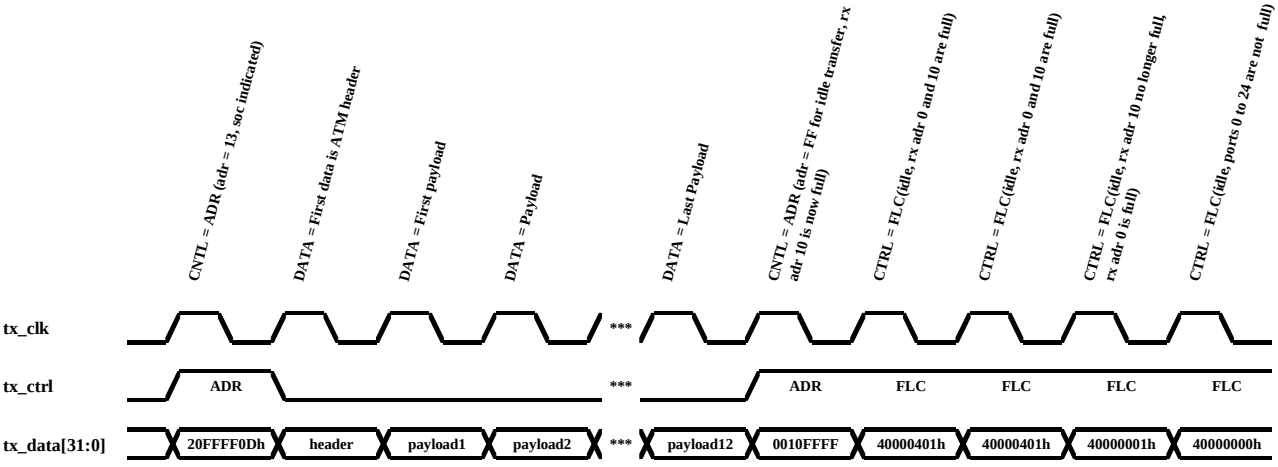

Figure 14: Idle Cell

Figure 15 demonstrates a packet transfer. It should be noted that in this example the value of N must be greater than 15 in compliance with Section 5.1.2, Item 3. Since transfer M is designated as the end of packet, the value of M may be less than 16.

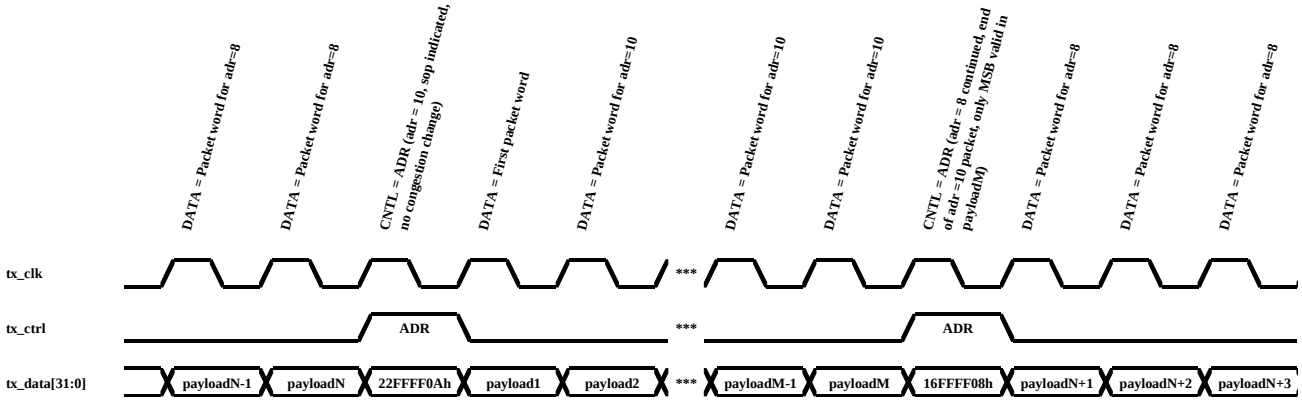

Figure 15: Packet Example

Figure 16 illustrates aborting a packet during mid-transmission. Since the packet is being deleted, J maybe less than 16.

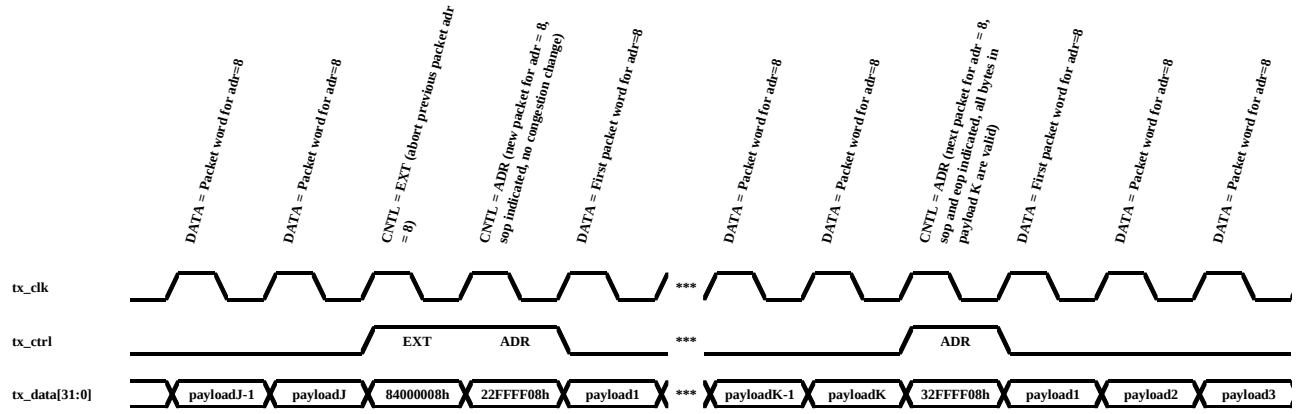

Figure 16: Aborting a Packet

9 I/O DEFINITION AND ELECTRICAL CHARACTERISTICS

The I/O technology is differential LVDS running at up to 415 MHz (T= approx. 2.41 nsec). This technology is capable of crossing a connector between a motherboard and daughter card. It is also capable of higher frequencies allowing for future expansion.

DC Characteristics

The IO characteristics shall conform to either ANSI/TIA/EIA-644 or IEEE STD1596.3–1996 standards. Table 7 is an informative summary of ANSI/TIA/EIA-644. Refer to these standards for more information.

Table 7: DC Specifications

Parameter	Description	Min	Max	Units
V_{OD}	Differential Output Voltage	247	454	mV
V_{OS}	Offset Voltage	1.125	1.375	V
Delta V_{OD}	|Change in V_{OD}|		50	|mV|
Delta V_{OS}	|Change in V_{OS}|		50	|mV|
I_{SC}	|Short Circuit Current|		24	|mA|
t_r/t_f	Output Rises/Fall Times	0.26	1.5	nsec
V_{TH}	|Threshold Voltage|		100	|mV|
V_{IN}	Input Voltage range	0	2.4	V

9.2 A.C characteristics

The AC characteristics are based on the timing specification for the receiver side of a signal. The setup and hold times are defined with regard to a positive clock edge as shown in Figure 17

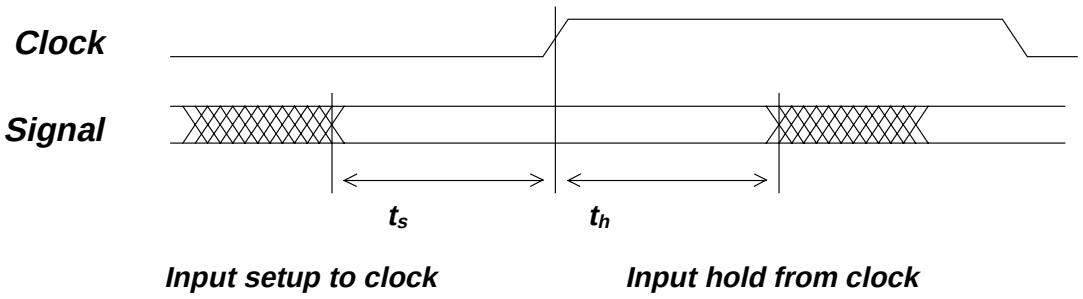

Figure 17: Setup and hold Time Definition

Table 8 below lists the AC characteristics for the Utopia Level 4 interface. The Test Conditions are the standard 100-ohm terminating resistor for LVDS with 25pF of capacitance. Utopia Level 4 is specified such that all inputs and outputs are registered. Thus, the inputs have setup and hold times

specified and the outputs have clock-to-output times specified.

Table 8: AC Specifications

Sym	Parameter	Test Condition	Min	Max	Unit
f_{max}	Clock Frequency			415	MHz
t_{PHL} t_{PLH}	Output Clock to DATA/CTRL Output	100 ohm differential, 25 pF each output	0.5	1.0	ns
t_s	Input Setup Time DATA/CTRL referenced to Input Clock	All Inputs	1.0		ns
t_h	Input Hold Time DATA/CTRL referenced to Input Clock	All Inputs	0		ns
duty	Clock Duty Cycle	High or Low	40	60	%

Note: When Clock is at the minimum period (approx. 2.41 nsec) the above device AC characteristics allows 0.2 nsec for board level clock pulse width distortion and 0.41 nsec for differential delay between Clock and DATA/CTRL.

10 Informative Appendix

The following sections are intended to aid in the understanding of this specification.

10.1 State Diagram Design Description

In order to aid in the design of the State Diagram shown in Table 9, the following Receiver State Transition Table is given:

Table 9: Protocol for Receive State Transition Table

Current State	CTRL	ADR	match	SOCP	EOP	Abort	Next State	Comments
IDLE	N	X	X	X	X	X	IDLE	Data xfer for other channel or overrun of last packet
	Y	N	X	X	X	X	IDLE	FLC, EXT or UDP
	Y	Y	N	X	X	X	IDLE	ADR for another channel
	Y	Y	Y	N	X	X	IDLE	New packet not starting yet
	Y	Y	Y	Y	X	X	START	Data for new packet is coming
START	N	X	X	X	X	X	ACTIVE	First data of new packet
	Y	N	X	X	X	N	START	FLC, EXT or UDP
	Y	Y	N	X	X	X	IDLE	Switched channels before data started - see design notes
	Y	Y	Y	N	N	X	START	A do nothing control, keep waiting for data
	Y	Y	N	N	Y	X	IDLE	Ended data before it started - see design notes
	Y	Y	Y	N	N	X	START	A do nothing control, keep waiting for data
	Y	N	Y	X	X	Y	IDLE	Abort
ACTIVE	N	X	X	X	X	X	ACTIVE	Data
	Y	N	X	X	X	N	ACTIVE	FLC, EXT or UDF
	Y	Y	N	X	N	X	PAUSE	Switching channels mid-packet
	Y	Y	N	X	Y	X	IDLE	End of packet
	Y	Y	Y	N	N	X	ACTIVE	"Do nothing" control
	52nd Data Byte Received						IDLE	End of cell (ATM only)
	Y	Y	Y	Y	N	X	START	Data aborted in mid-packet and start another
	Y	Y	Y	Y	Y	X	START	End of packet and start another
	Y	N	Y	X	X	Y	IDLE	Abort
PAUSE	N	X	X	X	X	X	PAUSE	Data for another channel
	Y	N	X	X	X	N	PAUSE	FLC, EXT, UDF
	Y	Y	N	X	X	X	PAUSE	Ctrl for another channel
	Y	Y	Y	N	X	X	ACTIVE	Resume transfer of data
	Y	Y	Y	Y	X	X	START	Abort paused packet and start another
	Y	N	Y	X	X	Y	IDLE	Abort paused packet

Design notes:

1. This design assumes that the receiving state machine described in Figure 7 and Table 9 does not inform its downstream logic about a new packet when it transitions into the START state, but waits until the first data is received, and it transitions to the ACTIVE state. Otherwise, the transitions back to IDLE from START are more difficult to manage.

2. This approach implies that the last byte of data in a packet cannot be sent downstream by the receiver until the CTRL transfer with EOP is sent indicating the number of bytes in it.

3. The sending state machine should probably be designed to wait until it has the beginning of a cell/packet before sending the CTRL transfer with its SOCP and should send the CTRL transfer with EOP right after the last data even if there isn't another packet to start. This seems like the simplest way to design and will result in the most efficient use of the protocol.

4. A cell/packet transfer may be interspersed with FLC, EXT and UDF transfers. One possible transfer is shown below in Table 10. In this example, There is an EXT transfer and an FLC transfer but no UDF transfer. If the receiver did not understand the UDF and EXT transfers, the data would still be correctly received.

   It will send the cell downstream when 52 bytes are received.

An ATM cell transfer may be interspersed with FLC, EXT and UDF transfers. One possible transfer is shown in Table 10 below. In this example, there is a UDF transfer and EXT transfer and an FLC transfer for each ATM cell. Although not defined in this proposal, the UDF transfer could put the ATM HEC in its correct position in the cell and the EXT could send parity bits over the whole transfer. At 17 clocks per cell a minimum clock rate of 400MHz is required to support full STS-192c (STM-64/AU-4-64c) bandwidth. If the receiver did not understand the UDF and EXT transfers, the ATM data would still be correctly received.

Table 10: Example of Data Transfer

tick 1	**tick 2**	**tick 3**	**tick 4**	**tick 5-16**	**tick 17**
FLC	ADR + SOCP	ATM DATA	EXT HEC	ATM DATA	EXT parity

tick 1	**tick 2**	**ticks 3-46**	**tick 47**	**tick 48**
FLC	ADR + SOCP	Packet DATA	ADR+EOP (may include next ADR+SOCP)	EXT parity

10.2 Addressing Example

For SONET/ Synchronous Digital Hierarchy (SDH) interfaces, the addressing shall be determined by the natural SONET port numbering associated with the lowest level SONET tributary supported by the PHY device. Examples of this numbering are shown in Table 11 - in all cases, it is assumed that all SONET interfaces are terminated by the same device that also contains the Utopia 4 interfaces. These examples also assume that the total potential capacity of the device is that of a SONET OC-192/ STM-64 stream, or equivalent.

Table 11: Correspondence between Ports & Addresses for Various SONET Configurations

Port Configuration Description Example	Port Numbering	Corresponding Addresses
Single OC-192 interface	1	0
One OC-192 signal containing four STS-48c payloads	1-4 (each STS-48c)	0-3
Four OC-48 signals, each carrying four STS-12c payloads	1-16 (1-4 for each STS-12c within the 1st STS-48, 5-8 for each STS-12c within the 2nd STS-48, etc.)	0-15
One OC-192 signal carrying 192 STS-1 payloads	1-192 in the standard SONET interleaving sequence (the sequence of port numbers in the data in the SONET stream is 1, 4, 7, …, 190, 2, 5, 8, …, 191, 3, 6, 9, …, 192)	0-191
Four OC-48 signals. The first, third and forth contain STS-48c payloads; the second contains 16 STS-3c payloads	1 (1st STS-48c payload), 17-32 (16 STS-3c payloads), 33 (STS-48c payload in 3rd OC-48), 49 (STS-48c payload in 4th OC-48). (Since the payload granularity over this interface is STS-3c, there are potentially 64 ports, in groups of 16)	0, 16-31, 32, 48

10.3 1-Bit Bus Width

In support of a simplex link, a low pin count mechanism for flow control can be accomplished by simply scaling the bus width to a single bit.

All DATA and CTRL word formats are as described in section 4.1, with a logical word size of 32-bits. Hence, a logical word transfer takes 32 clock cycles as shown in Figure 18. Any instance of the interface operates independently, with its own address space, flow control, clock and control pins. At the nominal clock rate, this interface can transfer up to 312 Mbps.

	Most Significant Byte [31:24]			Least Significant Byte [7:0]
Data1	A	B	C	D
Control	E	F	G	H
Data 2	I	J	K	L
Data 3	M	N	O	P

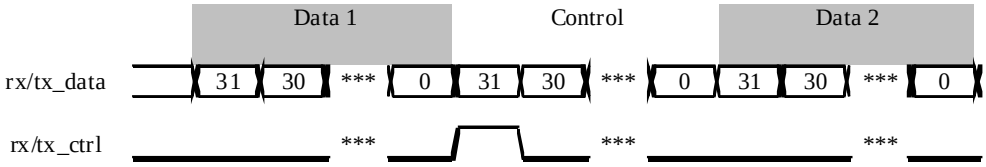

Figure 18: Transfer of 1-Bit Mode

The **rx/tx_ctrl** signal is set to '1' for the first clock cycle of the 32-bit logical control word, and is '0' otherwise. This permits unambiguous 32-bit word delineation on the 1-bit bus. Rx/Tx _data (31) is transmitted first, followed by rx/tx_data (30), until finally rx/tx _data (0) shown in Figure 18. If the data were an ATM cell, byte A would represent the GFC and the most significant bits of the VPI. Byte D would represent the least significant bits of the VCI, PTI and CLP bits.