


SpaceWire Evolutions



Trafe-Offs around Half Duplex SpW

***18th SpaceWire Working Group, 24th April 2012
Noordwijk, Netherlands***

Current Status:

- Networking technology for building on-board communications in S/C, used for the interconnection of:
 - Mass-memory
 - OBC
 - Telemetry
 - ...
- Designed by ESA and widely used on many ESA, NASA, JAXA, RKA space missions
- The standard specifies point-to-point full duplex links, with flow control mechanism which ensures that no data is lost due to receiver buffer overruns

The Problem:

- Bidirectional flow of data is not always required (e.g. sensors, actuators)
- One D-S pair is only used for FCTs:
 - Adds unnecessary mass since half of the wiring is practically unused
- Simplex cannot allow PnP, FDIR, operation in scheduled networks and may result in excessive data loss

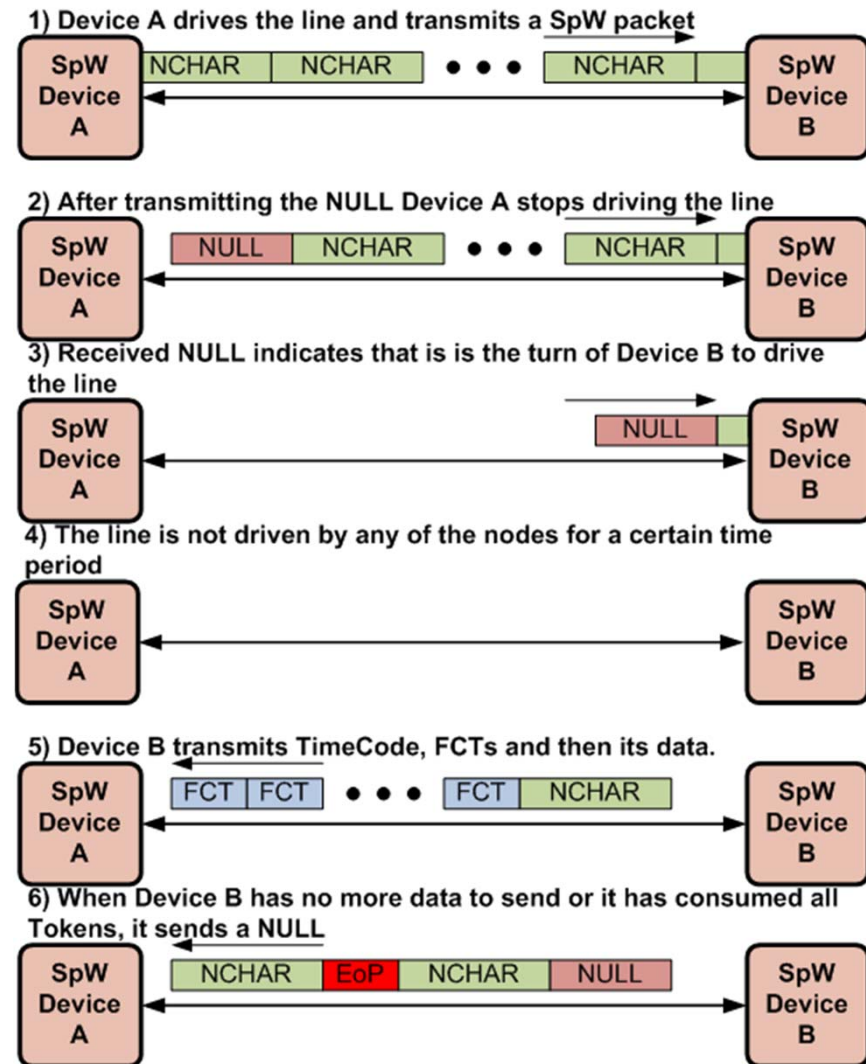
The proposed Solution:

- 4Links has proposed a solution for a Half Duplex version of SpW for asymmetric data transfers in which a single pair of D-S differential signals is shared between the two ends of the link
- The two ends alternatively act as transmitter and then receiver

Half Duplex SpW – 4Links Proposal

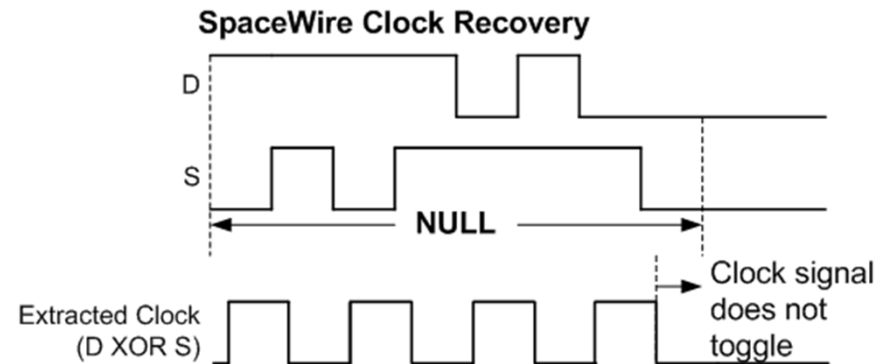
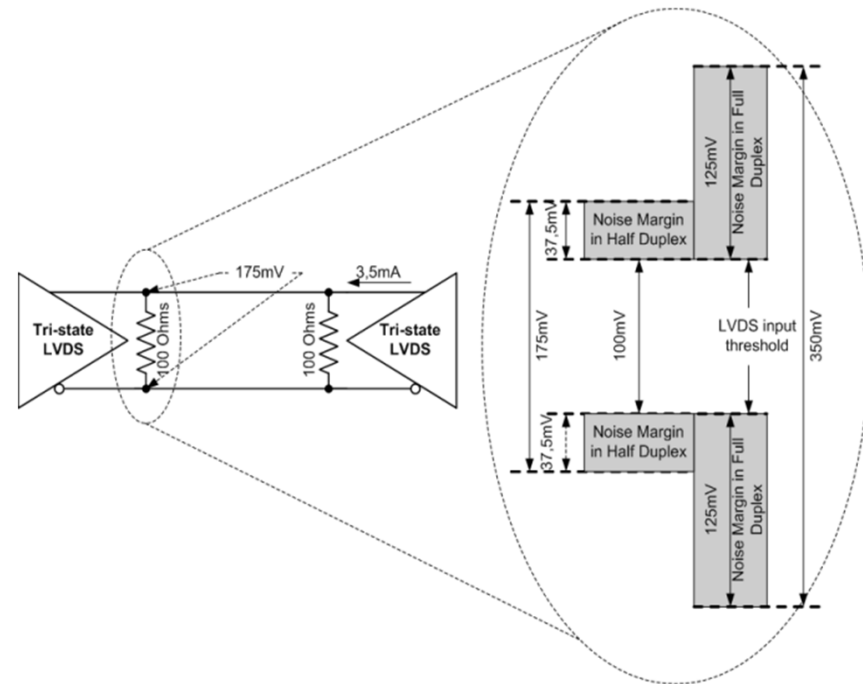
Half Duplex SpW:

- One end transmits its data until its buffer is empty, or
 - It has consumed all the FCTs it has received
 - It then sends a NULL indicating its has ceased transmission
 - Upon reception of a NULL the other end,
 - Sends a NULL if there is nothing to send, or
 - Transmits Time Codes, then
 - FCTs, then
 - NCHARs, EoP/EEP
 - And finally a NULL to enable the other end to resume transmission
- ⇒ **Half Duplex SpW offers all Full Duplex SpW features**



Half Duplex SpW – Technical Issues

- The main challenges with Half Duplex SpW are the Link Initialization, the link direction reversal and the Signal and Physical Levels (Rx pairs are the same as Tx pairs)
- Link Initialization:
 - Full Duplex state machine ensures that the two ends pass through the same states concurrently
 - Not possible with Half Duplex since they will be both listening or driving the line at the same time
- Direction reversal:
 - SpW Receivers extract the remote end transmission clock by XORing the D and S signals
 - After the last NULL is received no mode clock pulses are generated and the logic generating the “NULL_received” cannot be generated (logic remains unlocked)
- Half Duplex Signal Level:
 - LVDS is point-to-point and unidirectional
 - At some point in time both ends may be driving the line
 - Termination at both ends exceeds of the link causes the voltage at the termination resistors to be very close to the LVDS threshold
- Connector/cabling definition
 - One D-S pair wiring is not used



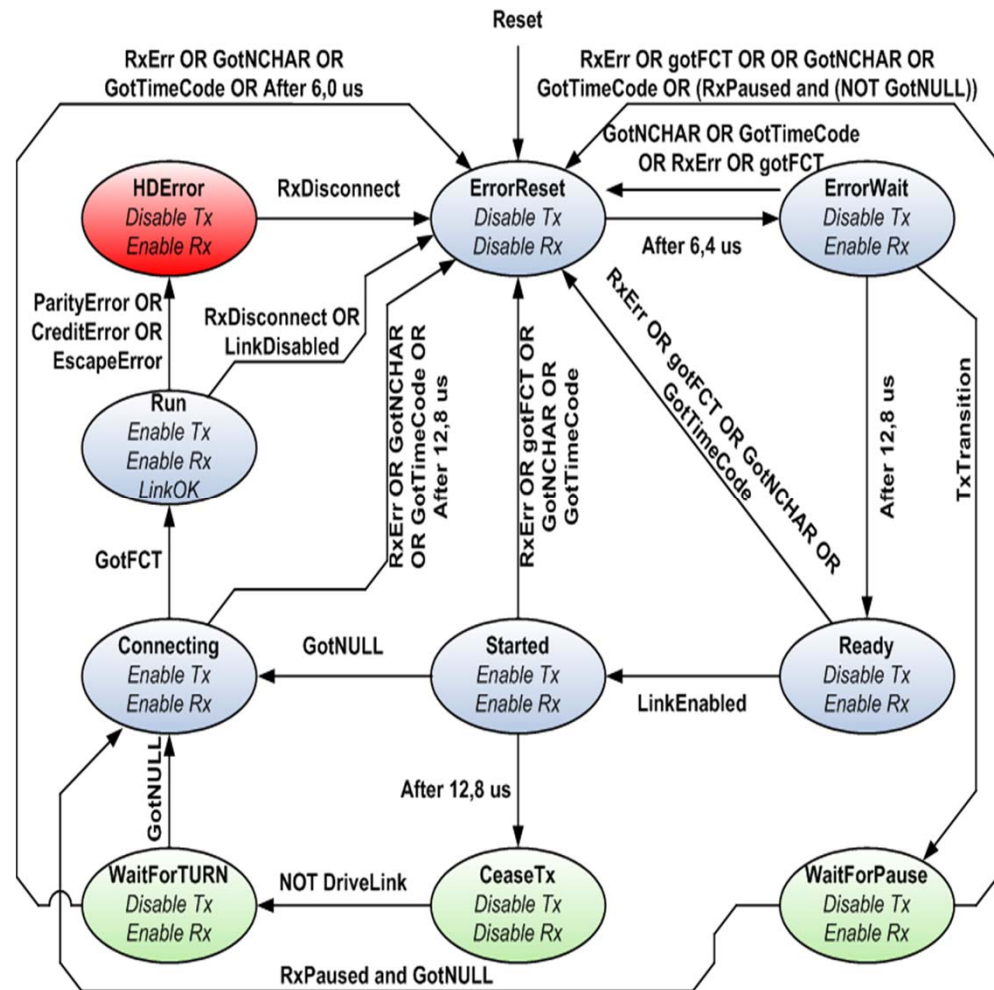
Half Duplex Link Initialization State Machine

New signals:

- **TxTransition:** Indicates that a transition has been detected on the D-S pair
- **RxPaused:** Indicates that no transition has been detected on the D-S pair for 200 ns
- **DriveLink:** Indicates that the link is driven by the controller's transmitter. Deasserted after the last character has been transmitted and the transmitter does not drive the link

New states:

- **WaitForPause:** Entered from the **ErrorWait** state if a transition has been detected on the transmit D-S pair. The state machine waits here until the remote end does not drive the line (rx_paused)
- **CeaseTx:** Entered from **Started** state if no NULL has been received for 12,8 us a condition which may indicate that the remote end is a full duplex and is disabled, or it is half duplex and waits its turn for transmission
- **WaitForTURN:** Entered from the **CeaseTx** state. At this state the transmitter is disabled to allow the remote end to transmit its NULLs and FCTs
- **HDError:** Entered when the receiving side detects a SpW Error. Ensures that the "Exchange of Silence" mechanism is followed upon SpW error occurrence




Half Duplex Link Initialization (1/3)

Controller 1 (Half Duplex):

- Is at state **Error Reset**
- A transition is detected on the D-S pair and the state machine proceeds to the **WaitForPause** state

Controller 2 (Full/Half Duplex):

- Is at state **ErrorReset**
- The 12,8 us expire and proceeds to the **Ready** and then **Started** state in which it transmits a NULLs
- Keeps sending NULLs at the **Started** state




Half Duplex Link Initialization (2/3)

Controller 1 (Half Duplex):

- After 200 ns the pause_rx signal is asserted and since it has got a NULL it proceeds to the **Connecting** state in which it transmits its FCTs
- Since it has not received FCTs from the remote side it transmits a NULL and returns to receive mode

Controller 2 (Full/Half Duplex):

- The 12,8 us interval expires and it proceeds to the **CeaseTx** state
- It ceases transmission and proceeds to the **WaitForTURN** state



Half Duplex Link Initialization (3/3)

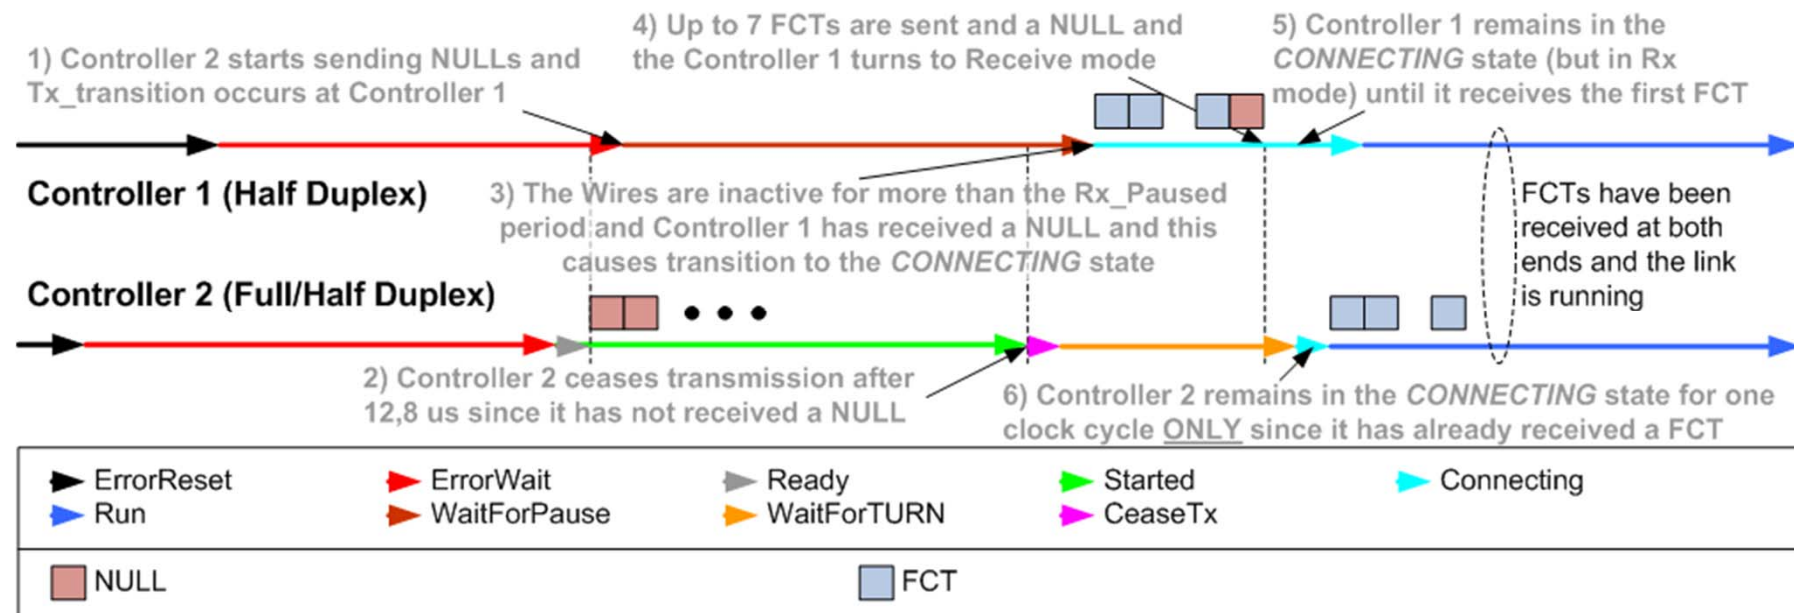
Controller 1 (Half Duplex):

- It receives the first FCT and proceeds to the **Run** state

Controller 2 (Full/Half Duplex):

- It receives the NULL and proceeds to the **Connecting** state and starts transmitting FCTs
- Since it has received FCTs from the remote side it proceeds to the **Run** state

Link is Initialized in Half Duplex mode



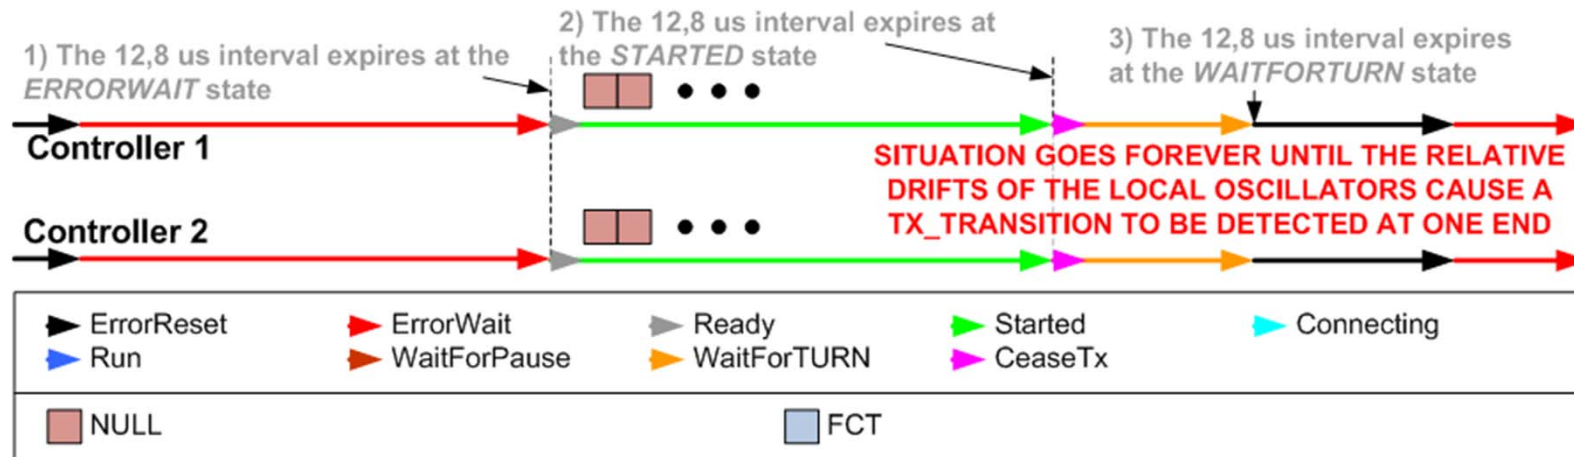
Half Duplex Link Initialization Livelock

The Problem:

- The two ends may be activated simultaneously
 - They will be passing from the same states at the same times
 - When one of them will be driving the link the other will do the same
 - When one of them will be listening the link the other will do the same
- This situation continues until the relative drifts of the local oscillator have cause enough drift for a tx_transition to be detected at one end
- If the links are clocked from the same source the link will never be initialized

Alternative Solutions:

- Add an offset to the 12,8 us timer (related to the port number) for devices with many links
- Pseudorandom offset at the 12,8 timer
- Positive/negative offset at the 12,8 us timer
 - Devices that have another link which is full duplex have positive offset (routers/concentrators)
 - Devices that do not have another link which is full duplex have negative offset (half duplex nodes)
- Configure nodes in auto-start mode and routers in Link Enabled



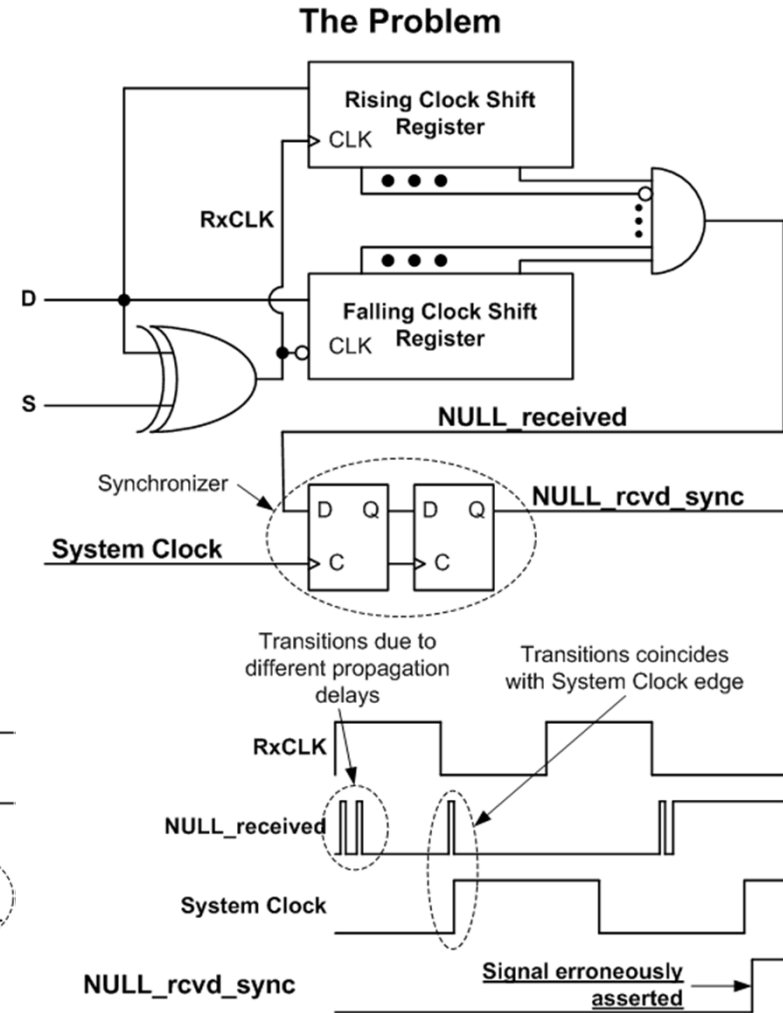
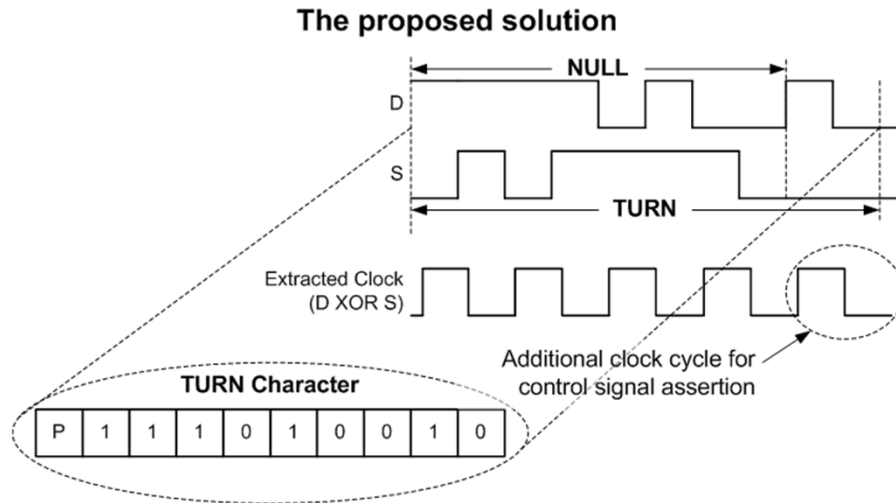
Link Direction Reversal

The Problem:

- The receiver's system clock and the reception clock are asynchronous
- The NULL_received signal will be generated by combinational logic
- The NULL decoder output has transitions due to changes of the logic levels at its inputs and differences in propagation delays
- These transitions may be patched by the system clock and erroneously cause link direction reversal

The proposed solution:

- Provide one more clock cycle to the receiver for decoder output latching by extending the NULL with a parity bit and one more zero (TURN character)



Half Duplex Signal Level (1/2)

Candidate Technologies			
Link Speed	LVDS (-A) (TIA/EIA 644 (-A))	M-LVDS (TIA/EIA 899)	BLVDS (not standardized)
Offset Voltage	1,125 – 1,375 V	0,3 – 2,1 V	1,185 – 1,435 V
Vout	454 mV (100 Ohms)	565 mV (50 Ohms)	350 (50 Ohms)
Transition time	260 ps	1000 ps	350 – 1000 ps
Driver strength	3,5 mA	11,3 mA	7 – 11,1 mA
Ground potential difference	±1 V	±2 V	±1 V
Input Voltage Range	0 – 2,4 V	-1,4 – 3,8 V	0 – 2,4 V
Input threshold	±100 mV	±50 mV	±100 mV
Max data rate (theoretical)	1,923 Gbps	500 Mbps	800 Mbps
Drivers contention	Not supported	Output current control	Output current control
Space Qualified Devices	Exist	Aeroflex UT54LVDM055LV ?	Aeroflex UT54LVDM031LV
Output voltage on 100 Ohms load	350 mV	1130 mV	700 mV – 1110 mV
Compatibility with LVDS	Yes	<u>Analysis per design is required. Current at LVDS termination resistor may cause a voltage of > 1 Vpp</u>	

Half Duplex Signal Level (2/2)

Tri-state LVDS


- LVDS drives 3,5 mA on the line
- 175 mV developed on the far end without taking losses into account which is above the LVDS input threshold, but marginally above
- Shorting two drivers on the line doubles the driving strength but worsens the eye pattern

B-LVDS:


- Designed for multi-drop topologies
- BLVDS drives 7 – 11 mA on the line
- 350 - 550 mV developed on the far end without taking losses into account
- Not an industry standard

M-LVDS


- Designed for multi-point topologies
- Input threshold is ± 50 mV around the CM voltage
- Industry standard (EIA/TIA-899)



National Semiconductors, Texas Instruments Recommendation




Xilinx Recommendation 1 (XAPP243)



Xilinx Recommendation 2 (XAPP232)

LVC MOS 25

Same Current on the cable as in Full Duplex SpW (3,5mA)



4Links recommendation

Higher current drive and turnaround affect EMC and cable definition

⇒ **Signal/Physical Level issues not covered within this study**

⇒ **Experimentation performed only for the prototype using SpW 1.0 connectors**

Half Duplex Signal Level – Initial Experiments

BLVDS

- Experiments with NI DS92LV010
- Cable Lengths 3, 10 meters

MLVDS

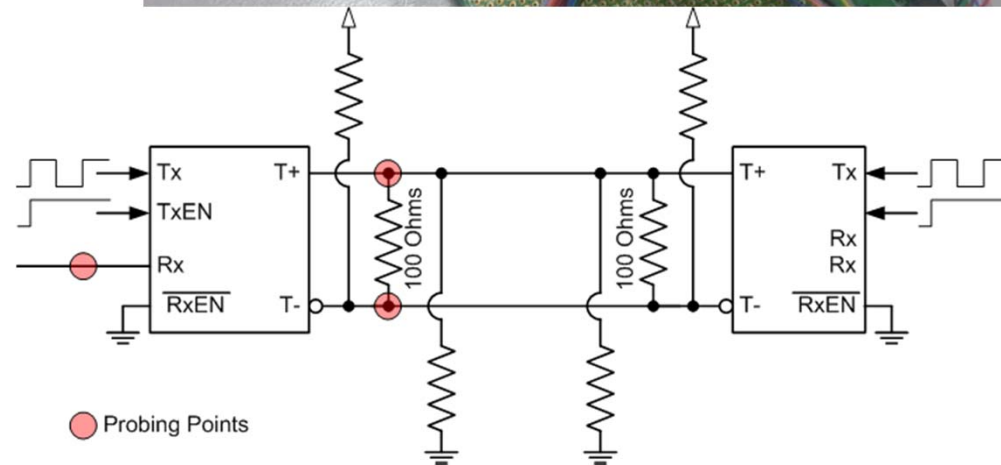

- Type 1 and Type 2 devices tested
- Experiments with SN65MLVD201 (Type 1) and with SN65MLVD206 (Type 2)
- Cable Lengths 3, 10 meters

Interoperability tests

- MVLDS - BLVDS
- MVDS – LVDS
- BLVDS – LVDS
- Pseudo B-LVDS – MLVDS & BLVDS

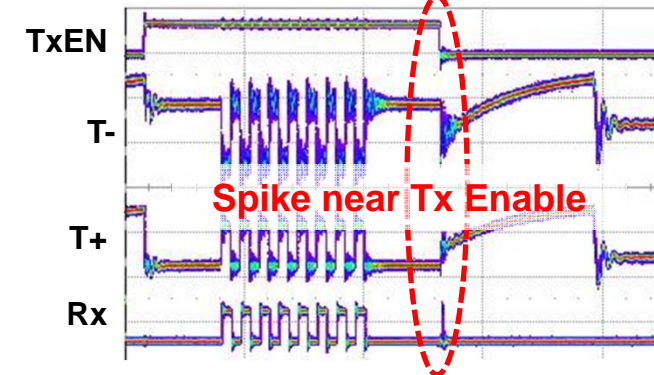



Probing points

- B/MLVDS differential signals
- B/MLVDS transceiver Rx LVTTTL output



Half Duplex Signal Level – Experimental Results

- Distortion at high speeds (Wire wrap mock up)
- BLVDS devices “prefer” parking the bus HIGH when not driven
- A spike appears near the TxEN signal edges
- Spike resolved by inverting the polarity of the D, S signals both at the transmitter and the receiver
- MLVDS and BLVDS present good interoperability




Half Duplex SpW & Latency (1/2)

Time Code (and packets) latency is increased

- A Time Code transmitter may not possess the bus
 - It shall wait the remote end to send a NULL
 - This time is proportional to the number of FCTs sent by the remote end, **plus**
 - The time the remote end needs to send its NCHARs, **plus**
 - The time it takes to send a NULL
- ⇒ **Half Duplex has excessive worst case Time Code and interrupts propagation latency**
- ⇒ **Increasing the FCTs increases efficiency but also increases, proportionally, the worst case latency**
- ⇒ **This worst case delay may occur per link**

Time Code worst case latency (us) vs. maximum number of FCTs						
Link Speed	Full Duplex	5 FCTs	7 FCTs	10 FCTs	15 FCTs	20 FCTs
10 Mbps	1	43,3	60,1	85,3	127,3	169,3
50 Mbps	0,2	9,06	12,42	17,46	25,86	34,26
100 Mbps	0,1	4,78	6,46	8,98	13,18	17,38
150 Mbps	0,067	3,35	4,47	6,15	8,95	11,75
200 Mbps	0,05	2,64	3,48	4,74	6,84	8,94
300 Mbps	0,033	1,92	2,48	3,32	4,72	6,12

Half Duplex SpW & Latency (2/2)



Available BW Utilization vs. FCTs (1/2)

The capacity of the receiver buffer affects Half Duplex SpW performance

- The transmitter transmits Nx8 NCHARs and then sends NULL waiting for FCTs to be received
- If the remote end has no data to send but only FCTs, this process inserts overhead time
- The overhead time consists of
 - 2 x turnaround time
 - 2 x 8(10) x bit time for NULLs
 - N x 4 x bit time for FCTs

- **Decreasing N:**
 - results in more frequent turnarounds
 - the turnaround and NULL times are not decreased

- ⇒ **Decreasing N decreases the Half Duplex SpW Performance**
- ⇒ **The impact of N is significant for high Link Speeds**
- ⇒ **Increasing N increases latency proportionally**
- ⇒ **N = 7 is a good compromise between efficiency, cost and latency**

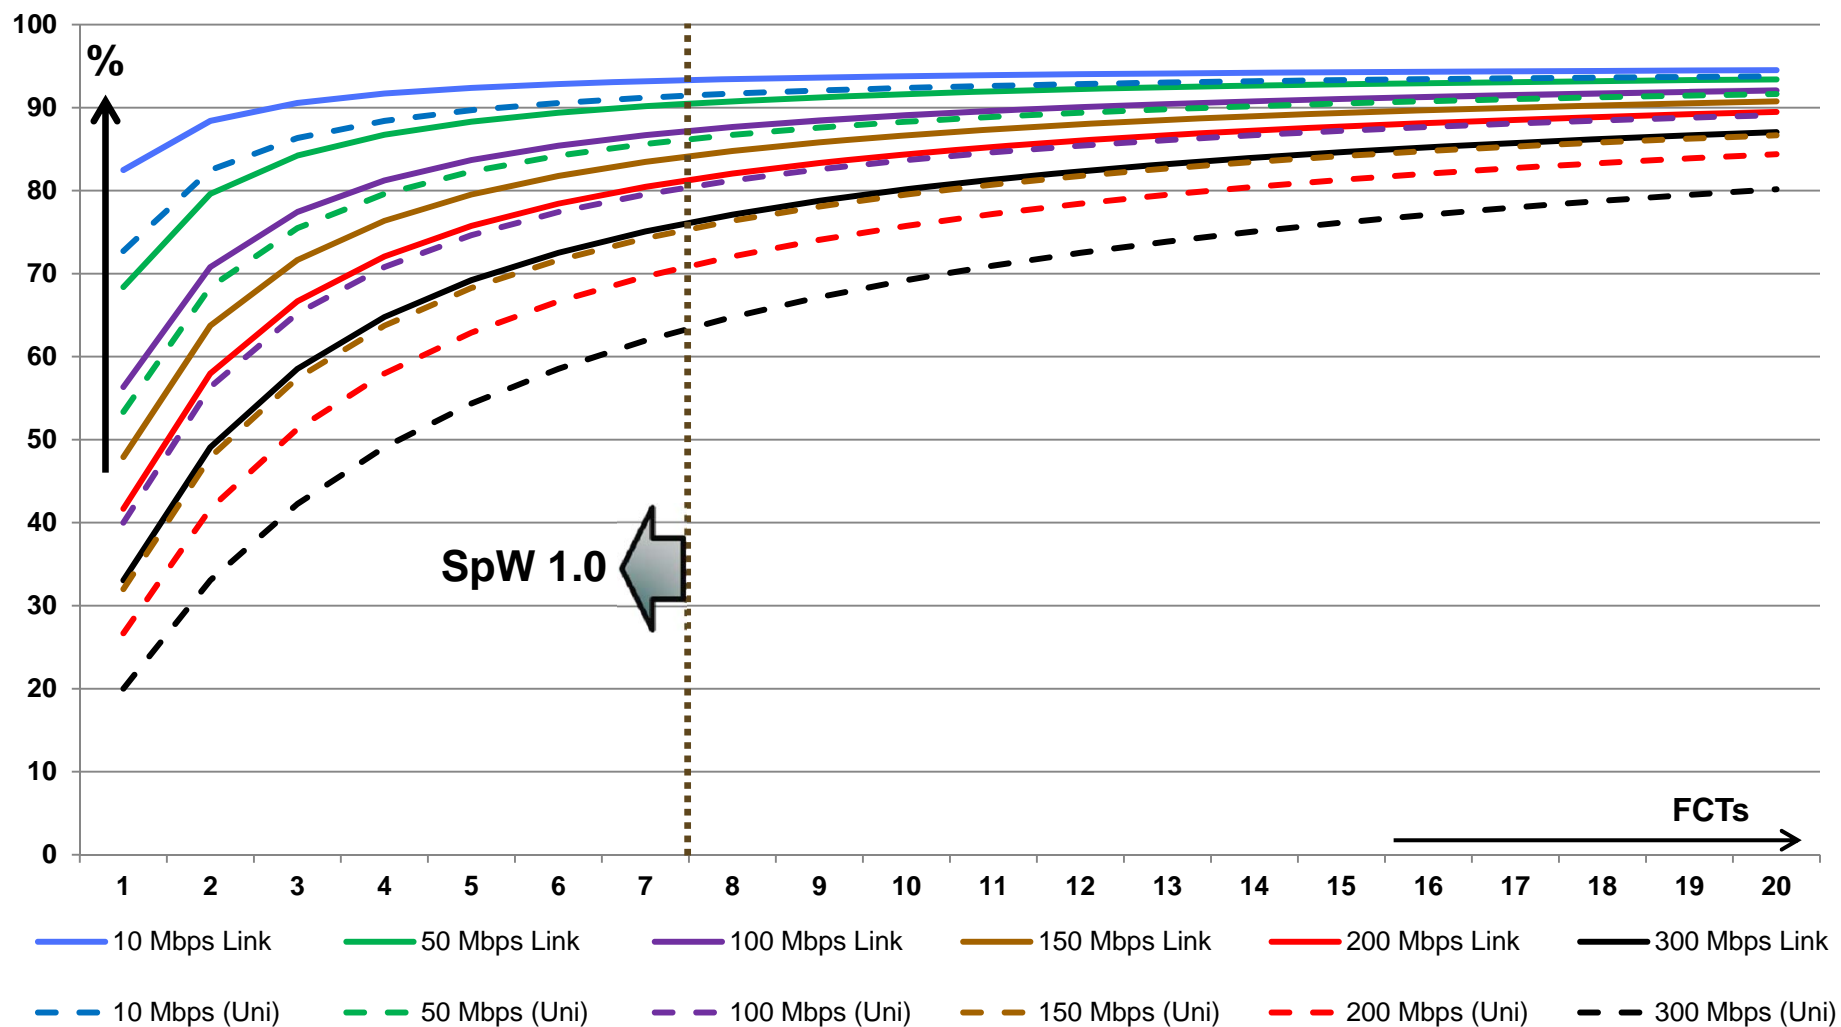
Available BW Utilization vs. FCTs for Unidirectional traffic

Link Speed	1 FCT	5 FCTs	10 FCTs	15 FCTs	20 FCTs
10 Mbps	73%	90%	92%	93%	94%
50 Mbps	53%	82%	88%	90.5%	91.5%
100 Mbps	40%	75%	83%	87%	89%
150 Mbps	32%	68%	79.5%	84%	86%
200 Mbps	27%	63%	76%	81%	84%
300 Mbps	20%	54%	69%	77%	80%


Available BW Utilization vs. FCTs for Bidirectional traffic

Link Speed	1 FCT	5 FCTs	10 FCTs	15 FCTs	20 FCTs
10 Mbps	82.5%	92%	94%	94%	94.5%
50 Mbps	68%	88%	91%	93%	93.5%
100 Mbps	56%	83%	89%	91%	92%
150 Mbps	48%	79%	86.5%	89%	90.5%
200 Mbps	42%	76%	84%	87%	89%
300 Mbps	33%	69%	80%	85%	87%

Available BW Utilization vs. FCTs (2/2)



FCTs impact on Available BW utilization



Available BW Utilization vs. turnaround time (1/2)

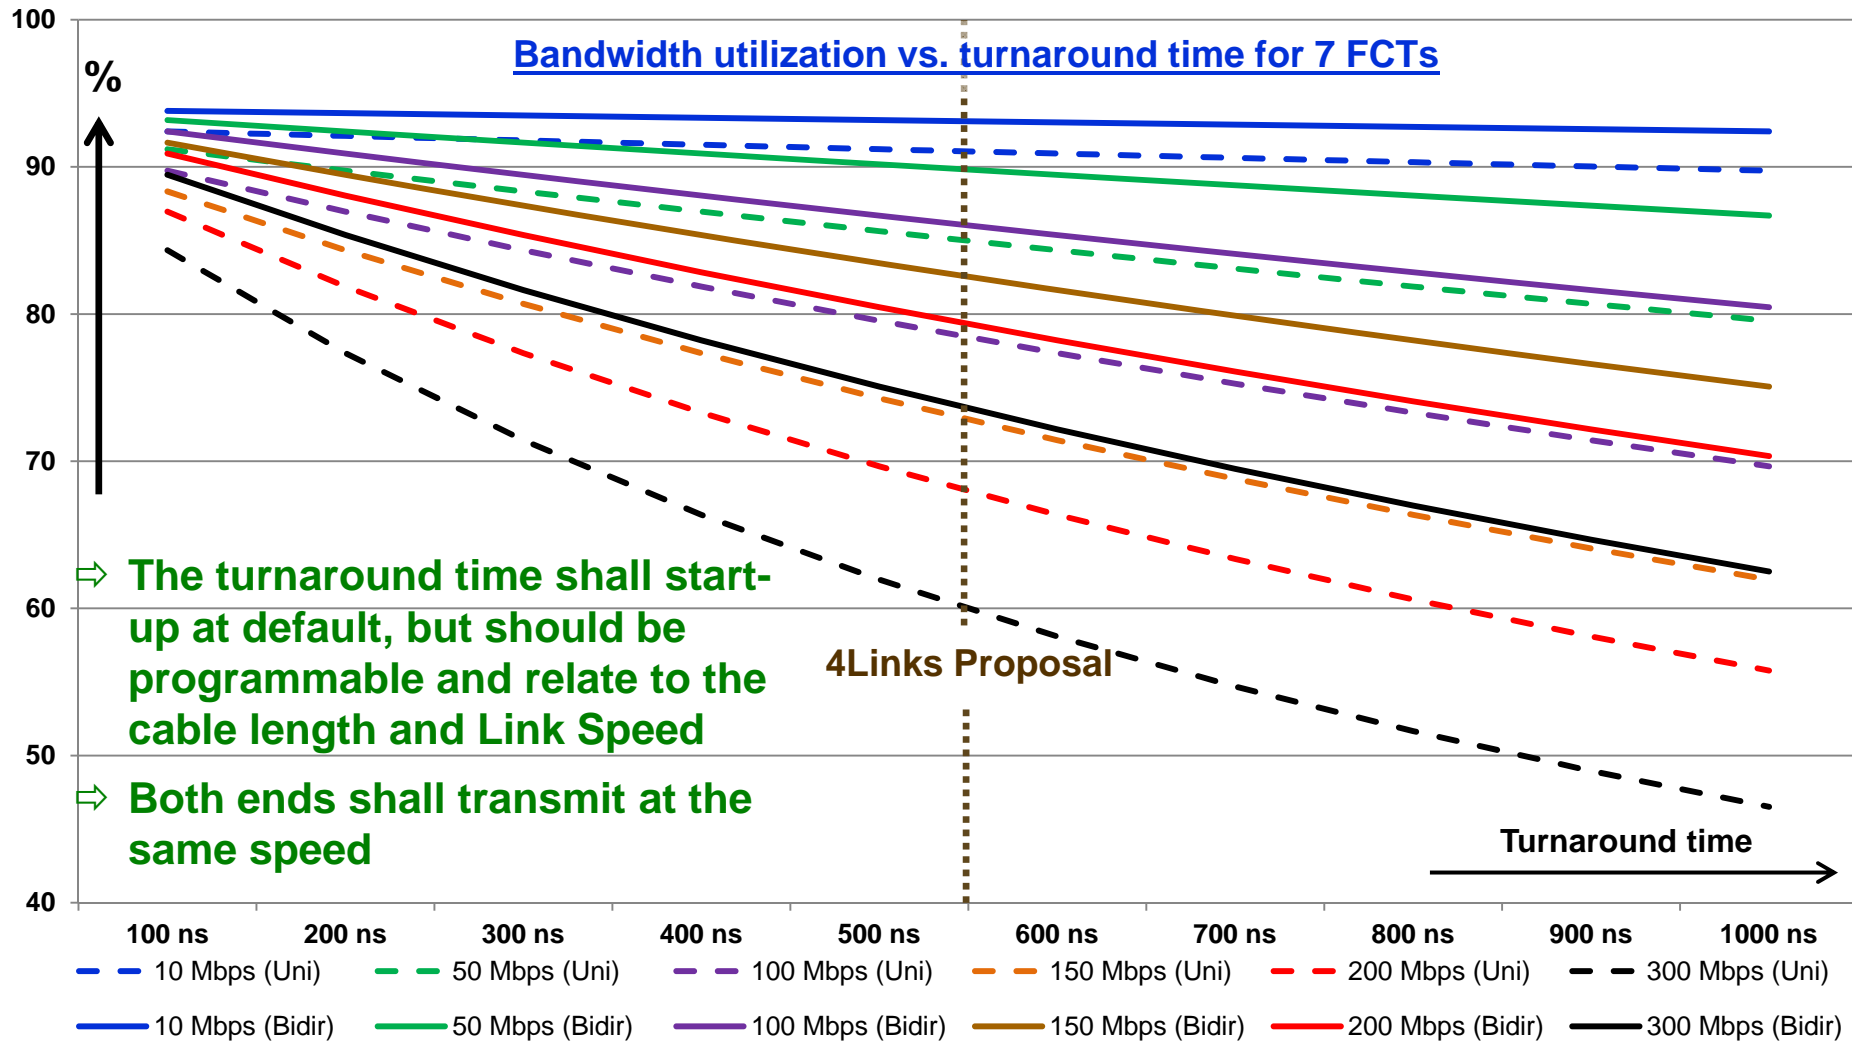
The turnaround time (time in which the link remains undriven) affects Half Duplex SpW performance

- The turnaround time consumes Bandwidth
 - The consumed Bandwidth increases as the turnaround time increases
 - The consumed bandwidth increases as the number of turnaround times per second increases
- ⇒ The turnaround time has a significant impact at high speeds
- ⇒ The bandwidth consumed by the turnaround is inversely proportional to the maximum number of FCTs sent from the receiver to the transmitter

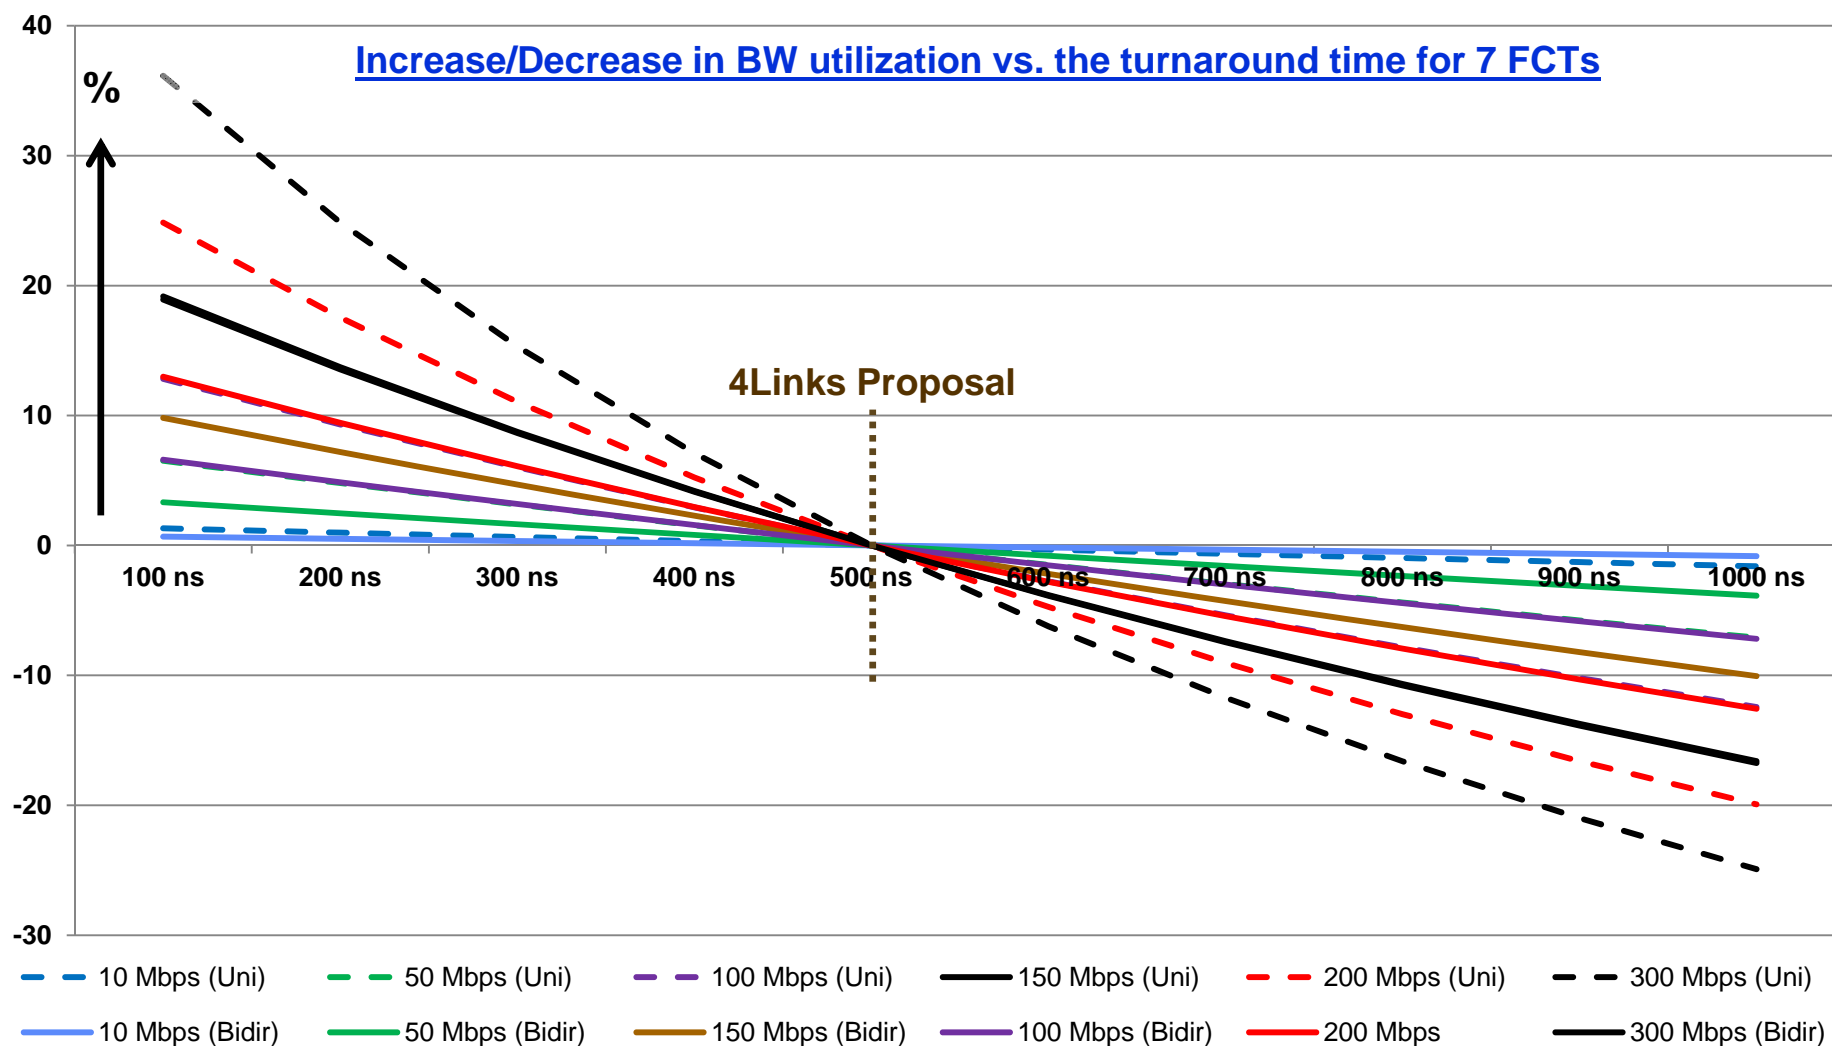
Available BW Utilization vs. Turnaround time for <u>Unidirectional</u> traffic					
Link Speed	200 ns	400 ns	600 ns	800 ns	1000 ns
10 Mbps	92%	91%	91%	90%	90%
50 Mbps	91%	87%	84%	82%	79%
100 Mbps	87%	82%	77%	73%	69%
150 Mbps	84%	77%	71%	66%	62%
200 Mbps	81%	73%	66%	60%	56%
300 Mbps	77%	66%	58%	51%	46%

Available BW Utilization vs. Turnaround time for <u>Bidirectional</u> traffic					
Link Speed	200 ns	400 ns	600 ns	800 ns	1000 ns
10 Mbps	93%	93%	93%	92%	92%
50 Mbps	92%	91%	90%	88%	87%
100 Mbps	91%	88%	85%	83%	80%
150 Mbps	89%	85%	81%	78%	75%
200 Mbps	88%	83%	78%	74%	70%
300 Mbps	85%	78%	72%	67%	62%


Available BW Utilization vs. turnaround time (2/2)



Turnaround time impact on Available BW utilization



FCTs impact on BW utilization vs. turnaround time



Conclusions

😊 Advantages:

- Supports all SpW 1.0 functionality and does not infer the hazard of NCHAR loss as Simplex SpW
- Wormhole routing supported
- Fair bandwidth allocation between the two ends of the link
- Requires simple functional changes in the SpW Cores logic since the functionality is almost identical
- Simpler and lighter cabling required – lighter to be confirmed after EMC characterization
- Lower cost solution for networks with few hops without inferring large jitter/latencies
- Proposed state machine allows for auto-detection of Full/Half Duplex

☹ Drawbacks:

- Physical Level modifications required – define connector to achieve optimized throughput/mass performance
- Signal Level modifications required
- BLVDS and M-LVDS inject 50% more current than LVDS and therefore cable definition shall be re-examined for EMC issues – cannot yet evaluate throughput vs. mass performance
- Character Level modification for the new “TURN” character is required
- Latency and Jitter is introduced in Time-Code propagation and application packets – not suitable as backbone network in scheduled networks with complex topologies
- Cannot support precise time-distribution
- May present excessive jitter in hot redundant topologies
- Efficiency and Latency are factors driving to opposite directions. Trade-off analysis per application is required

Fields of application:

- Concentrators which receive data asynchronously and propagate them through full Duplex SpW