Comparison of Channel Protocols for Fast, Low Energy Communication over Transmission Lines

Shomit Das*, Kenneth S. Stevens

The University of Utah

*now with AMD Research

Exascale Challenges

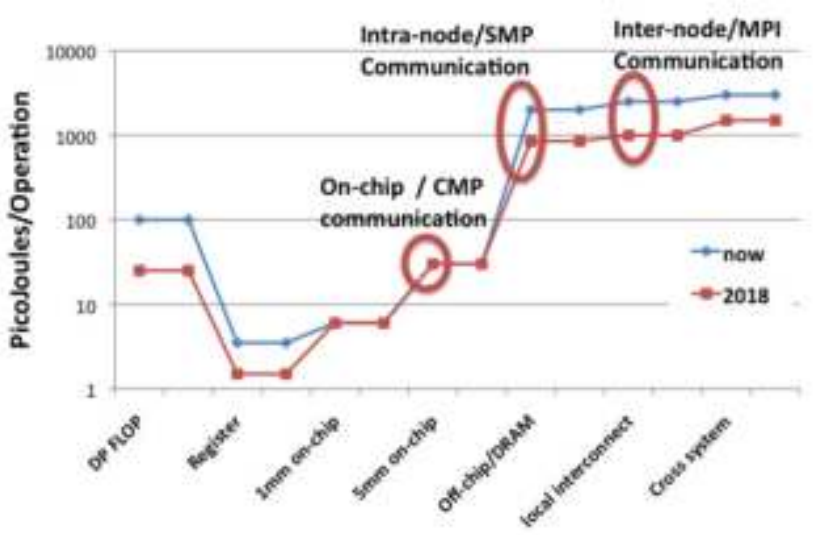

Cost of data movement relative to cost of a flop*

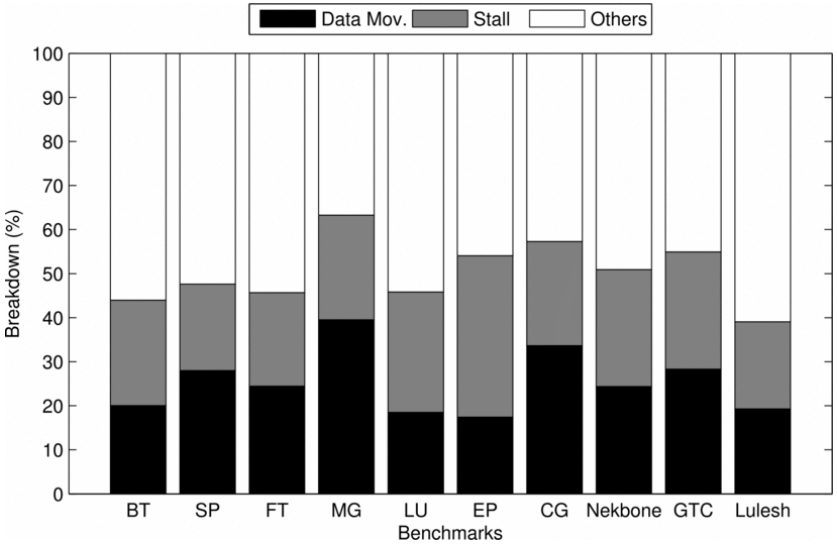

Data movement energy component**

* J. Shalf et.al., Exascale Computing Technology Challenges, LNCS 2011

** G. Kestor et.al., Quantifying the energy cost of data movement in scientific applications, IISWC 2013

Interconnect Challenges

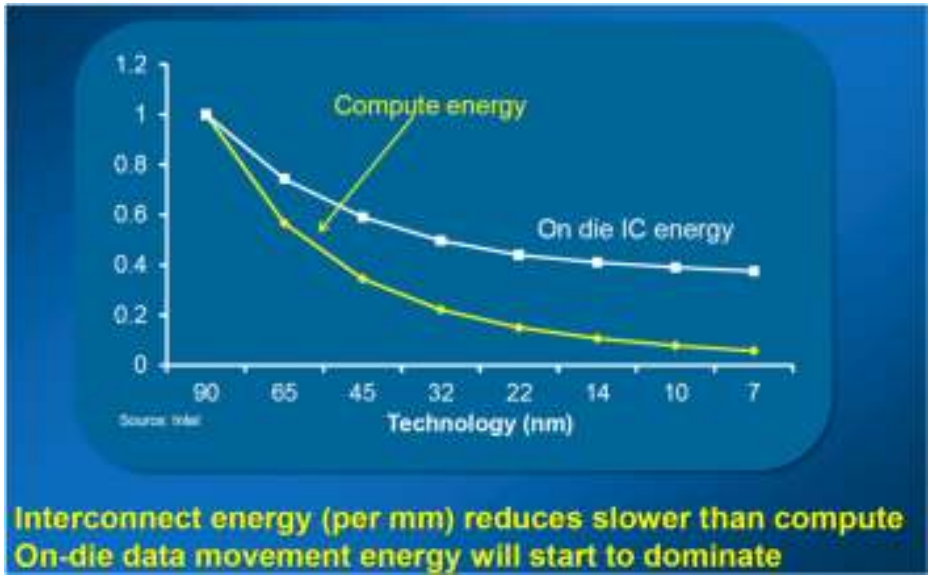

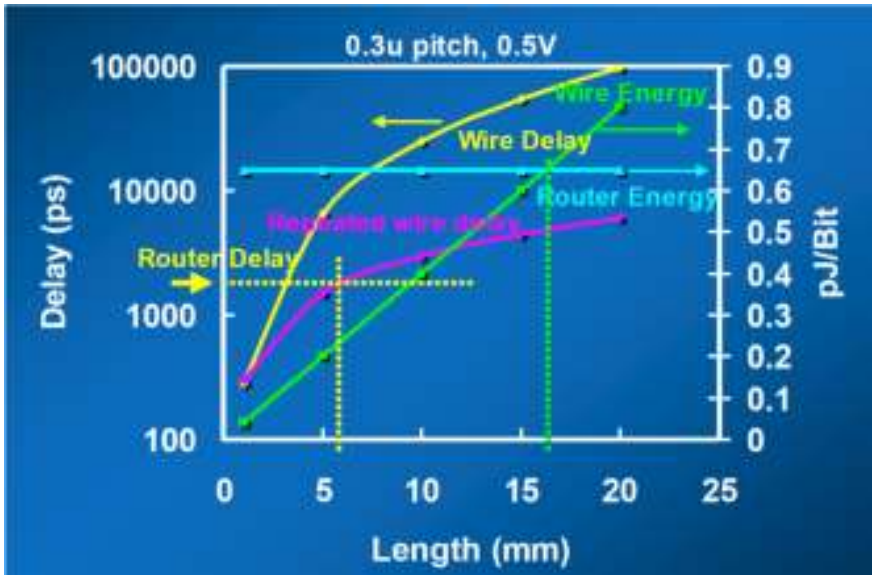

*Shekhar Borkar, Exascale Computing- Fact or Fiction? IPDPS 2013

Global Interconnect

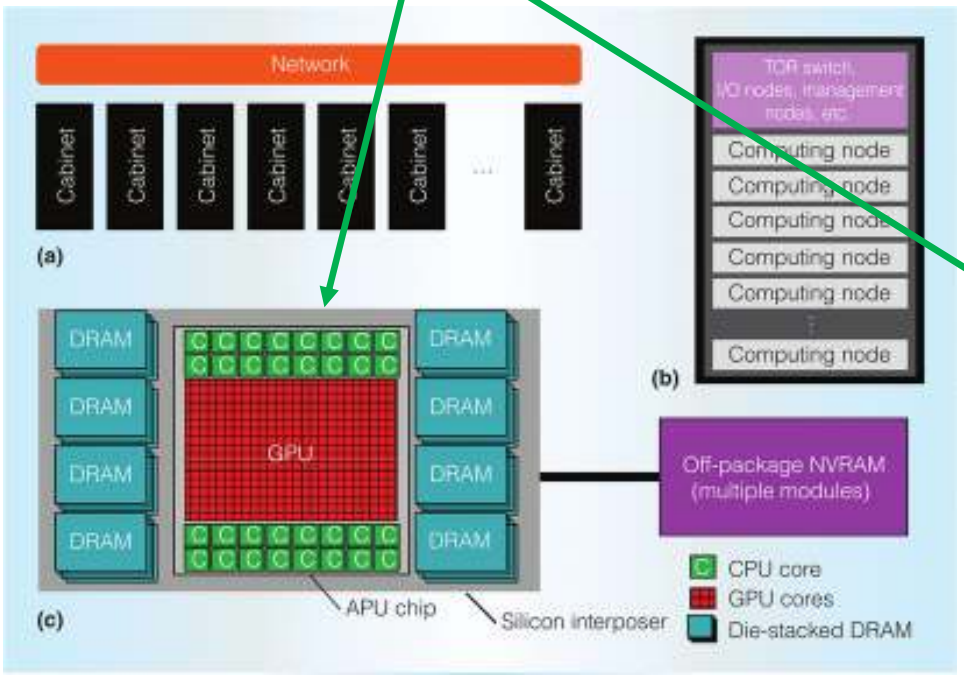

AMD

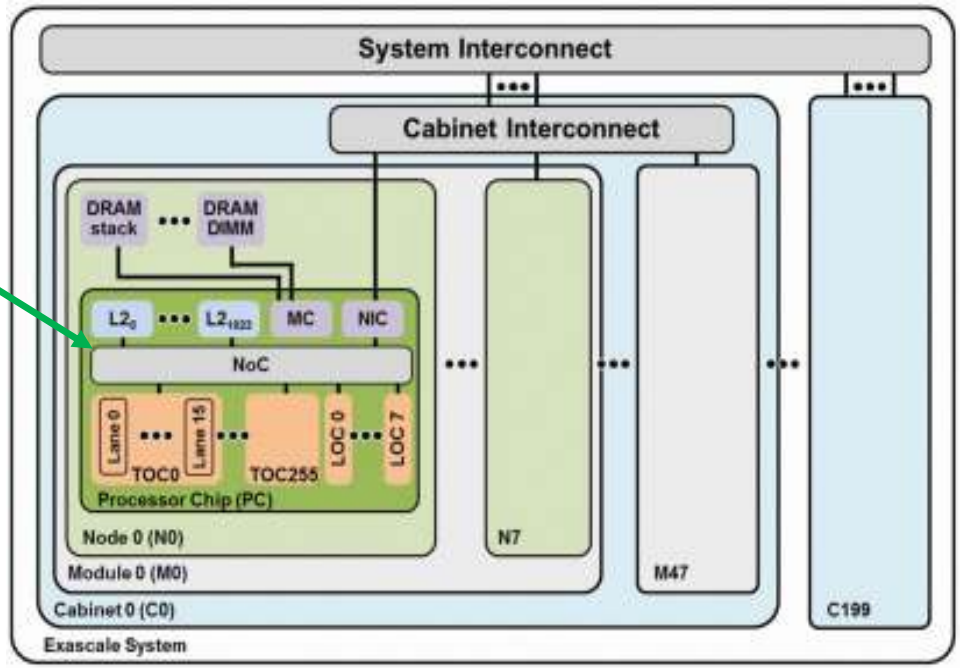

NVidia

Exascale System Architecture Examples (proposed)

Transmission Lines

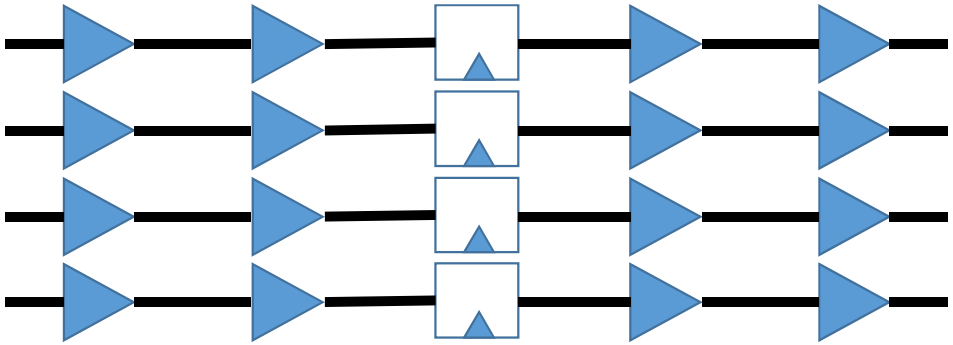

Repeated RC wire

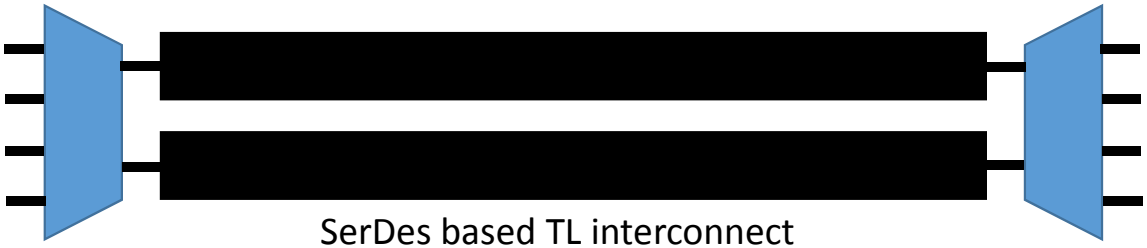

SerDes based TL interconnect

Transmission Lines require thick top level metals
They require carefully designed signal and return paths
Signal integrity depends on interconnect aspect ratio among many other factors
Bandwidth per unit area suffers as a result
Analog signaling techniques such as differential signaling, current mode signaling are applied
Higher frequencies can be used
SerDes means more timing and energy considerations

On-chip Transmission Lines

Transmission Lines – Design and Simulation

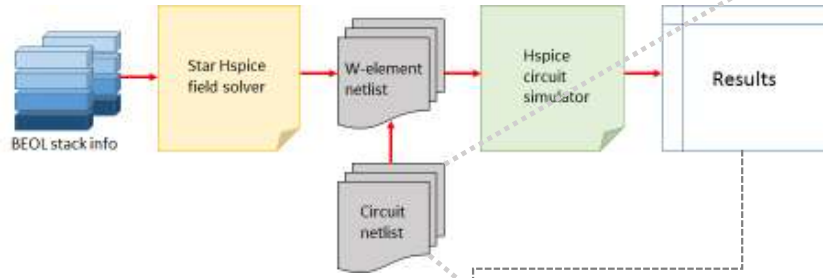

Transmission Line (opt)	Latency (ps)	Energy (pJ/bit)	Pitch (μm)	Eye Height (V)	Eye Width (ps)
TL1 (energy)	338.3	0.54	2.8	0.88	339.6
TL2 (latency)	292.9	0.60	7.6	0.85	337.9

7mm TL

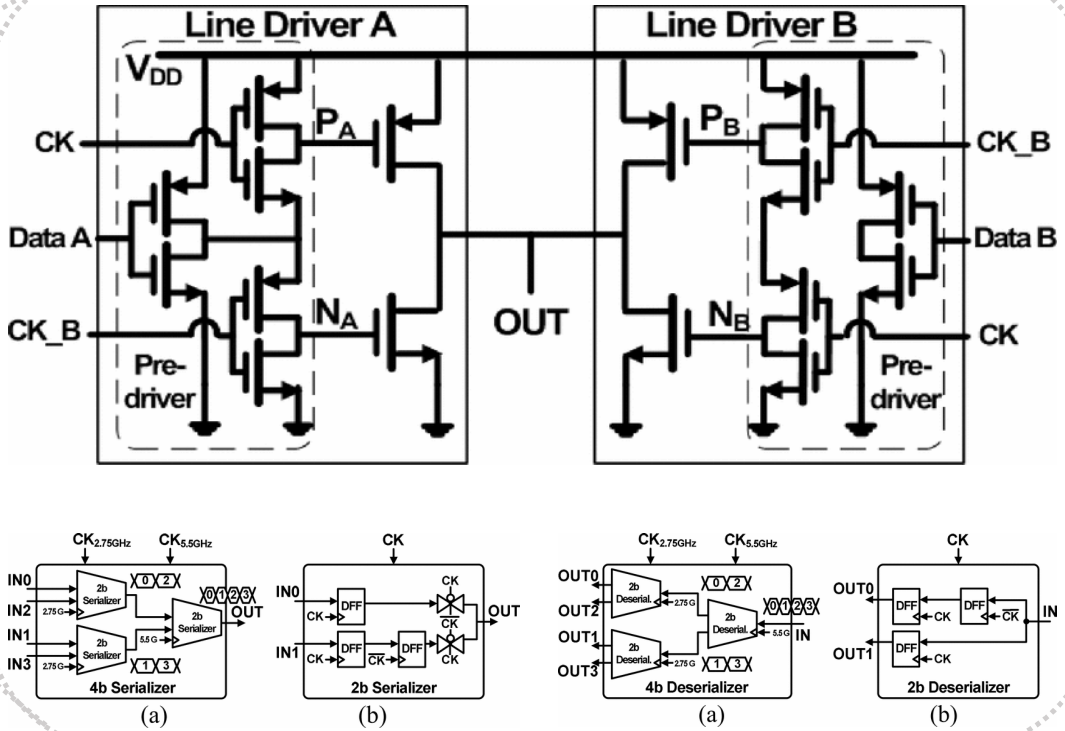

Transmission Line Interconnect Design Environment

*H.G. Rhew et.al., A 22Gb/s, 10 mm on-chip serial link over lossy Transmission Line, ESSCIRC 2012

Communication Protocols

- Dual Rail
- Bundled Data 4-phase
- Bundled Data 2-phase
- Source Asynchronous Signaling
- Clocked latched
- Clocked flopped
- Source Synchronous

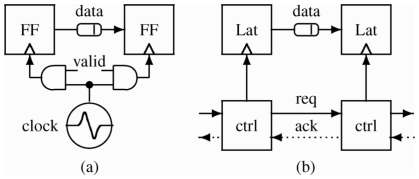

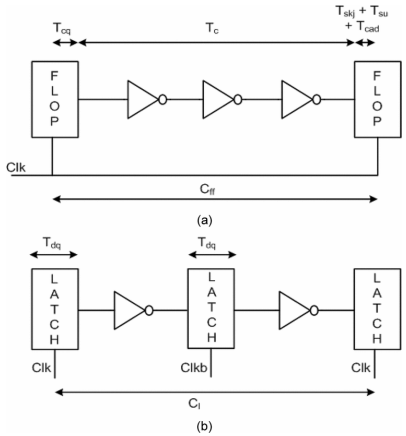

Communication Protocols - SAS

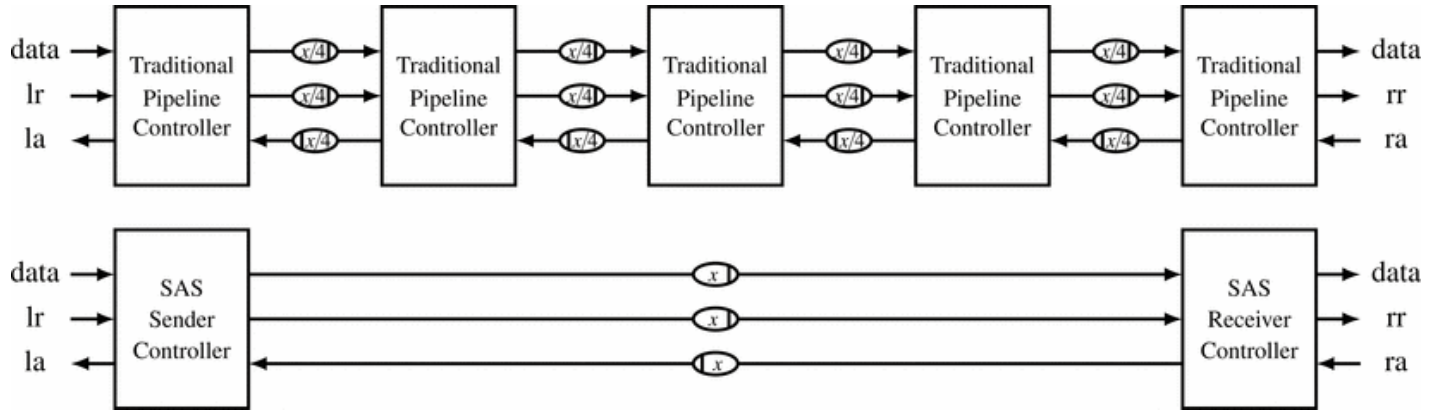

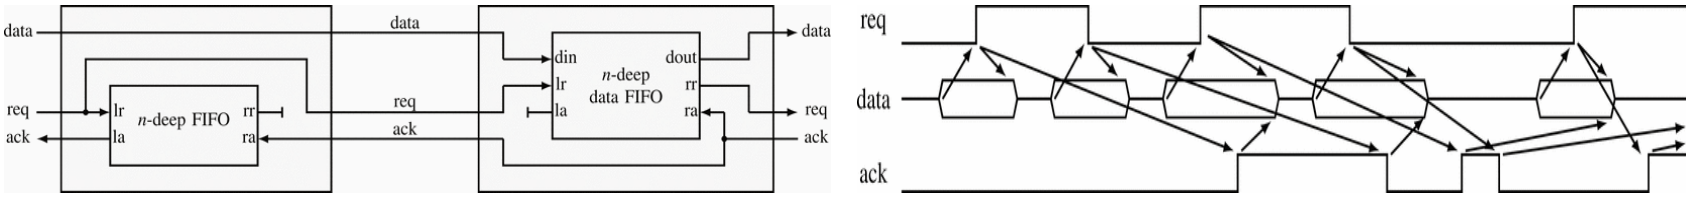

Source Asynchronous Signaling (uncoupling req and ack)

Metrics

Name	Description	FO4 Value
V_c	Cross coupling variation of wire	0.4
V_p	Process variation	0.16
V_{vt}	Voltage & temperature variation	0.12
T_{cq}	Delay from clock to output	1.8
T_{sul}	Latch setup time	0.48
T_{suf}	Flop setup time	0.96
T_{skj}	Clock skew and jitter	1
T_h	Hold time	0.25
T_{dqf}	Flop data-to-output delay	2.8
T_{dql}	Latch data-to-output delay	1.4
T_{pw}	Minimum pulse width that can be propagated through a critical distance	2.2
T_c	TL Channel Delay	5.0 – 6.3
T_{c+}	$T_c(1+\frac{V_c}{4})$	5.5 – 7.0
T_{c-}	$T_c(1-\frac{V_c}{4})$	4.5 – 5.7
T_{fsm+}	Controller delay per phase	2.55 – 6.64
T_{fsm-}	$T_{fsm+}(1-\frac{V_c}{4})$	2.34 – 6.10
T_{lat+}	Forward latency of async controller	1.18 – 3.95
T_{cfifo+}	Cycle time of 0-bit tree FIFO	12.12
$T_{latfifo+}$	Forward latency of 32-bit tree FIFO	4.33
E_{ctl}	Controller energy per phase per bit	0.072 – 0.90
E_{tl}	Energy of a TL link	1.32 – 2.12
E_{serdes}	Energy of a SERDES apparatus	2.15
E_{datsas}	Energy of a SAS data link	1.42
E_{ctlsas}	Energy of a SAS control logic	5.0
E_{rep}	Energy of a single repeater	1
E_{datf}	Data energy through flop	2.02
E_{datl}	Data energy through latch	1.78
E_{clkf}	Clock energy for flop	0.22
E_{clkl}	Clock energy for latch	0.2
E_{clkdef}	Clock energy for DETFF	0.42

Models – Cycle Time

Name		equation	description
T_{ps}	$\leq$	$(T_{fsm+} + T_{cq+} + T_{c+} + T_{su}) - (T_{fsm+} + T_{c-} + T_{fsm-})$	Handshake to data margin
T_{cdel}	$\leq$	$(1 + \frac{V_p}{2})\max(0, T_{ps}) - (1 - \frac{V_p}{2})\max(0, T_{ps})$	Robust delay element size
C_{ff+}	$\geq$	$\max(2T_{pw}, (T_{c+} + T_{su} + T_{skj} + T_{cq+}))$	Clk_flop throughput
C_{l+}	$\geq$	$\max(2T_{pw}, T_{c+} + 2T_{dql})$	Clk_latch throughput
C_{di+}	$\leq$	$4T_{fsm+} + 2T_{c+} + 2O_{wd}T_{c+}$	2-rail and 1-of-4 DI
C_{bd2+}	$\leq$	$2T_{fsm+} + T_{c+} + T_{cdel} + O_{wd}T_{c+}$	2-phase bundled data cycle time
C_{bd4+}	$\leq$	$4T_{fsm+} + 2T_{c+} + T_{cdel} + 2O_{wd}T_{c+}$	4-phase bundled data cycle time
C_{ss+}	$\leq$	$\max(2T_{pw}, (T_{cq+} + T_{c+} + T_{su}), (T_{c+} + T_{fsm+}))$	src-sync cycle time
C_{sas+}	$\leq$	T_{cfifo+}	SAS

Cycle Time expressions

Models – Latency

Name		equation	description
L_{ff+}	$\geq$	$\max(2T_{pw}, (T_{c+} + T_{su} + T_{skj} + T_{cq+}))$	Clk_flop latency
L_{l+}	$\geq$	$\max(2T_{pw}, T_{c+} + 2T_{dql})$	Clk_latch latency
L_{di+}	$\leq$	$T_{lat+} + T_{c+}$	2-rail and 1-of-4 DI latency
L_{bd2+}	$\leq$	$T_{lat+} + T_{c+} + T_{cdel}$	2-phase bundled data latency
L_{bd4+}	$\leq$	$T_{lat+} + T_{c+} + T_{cdel}$	4-phase bundled data latency
L_{ss+}	$\leq$	$\max(2T_{pw}, (T_{cq+} + T_{c+} + T_{su}), T_{c+} + T_{fsm+})$	src-sync latency
L_{sas+}	$\leq$	$T_{latfifo+} + T_{c+}$	SAS latency

Latency expressions

Models – Energy

Name	equation	description
E_{gf}	$\frac{1}{A_b}\frac{E_{clkf}}{12}+\frac{E_{clkf}}{4}$	Flop clock gate energy
E_{gl}	$\frac{1}{A_b}\frac{E_{clkl}}{12}+\frac{E_{clkl}}{4}$	Latch clock gate energy
E_{fe}	$\frac{1}{A_b}2P_{dc}E_{clktree}+\frac{W_b}{4}(2(E_{gf}+E_{clkf})+A_w(E_{datf}+E_{tl}+E_{serdes}))$	Clk_flop energy
E_{le}	$\frac{1}{A_b}2P_{dc}E_{clktree}+\frac{W_b}{4}(4(E_{gl}+E_{clkl})+A_w(2E_{datl}+E_{tl}+E_{serdes}))$	Clk_latch energy
E_{4e}	$4(E_{ctl}+E_{tl})$ $+\frac{W_b}{4}(2E_{clkl}+A_w(E_{datl}+E_{tl}+E_{serdes}))$	4-phase bundled data energy
E_{2e}	$2(E_{ctl}+E_{tl})$ $+\frac{W_b}{4}(E_{clkdef}+A_w(E_{datl}+E_{tl}+E_{serdes}))$	2-phase bundled data energy
E_{ss}	$2(E_{ctl}+E_{tl})$ $+\frac{W_b}{4}(2E_{clkf}+A_w(E_{datl}+E_{tl}+E_{serdes}))$	src_sync energy
E_{di2}	$\frac{W_b}{4}(4E_{ctl}+2E_{tl})+2E_{tl}+2E_{serdes}$	DI 1-of-2 energy
E_{sas}	$(E_{ctlsas}+E_{tl}+\frac{W_b}{4}(E_{ctlsas}+A_w(E_{tl}+E_{serdes})))$	SAS energy

Energy per transaction expressions

Comparisons – Cycle Time

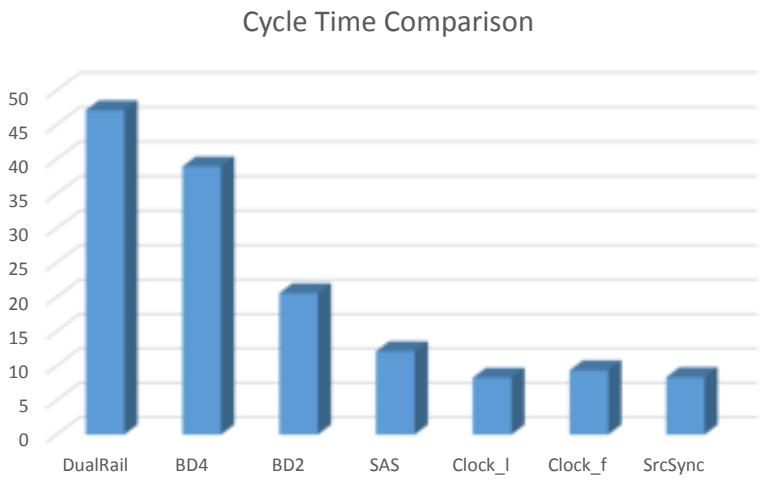

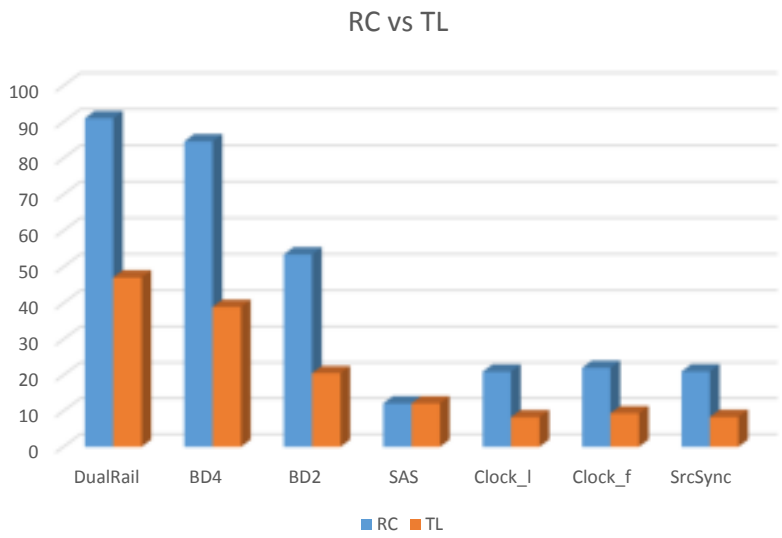

Comparisons - Latency

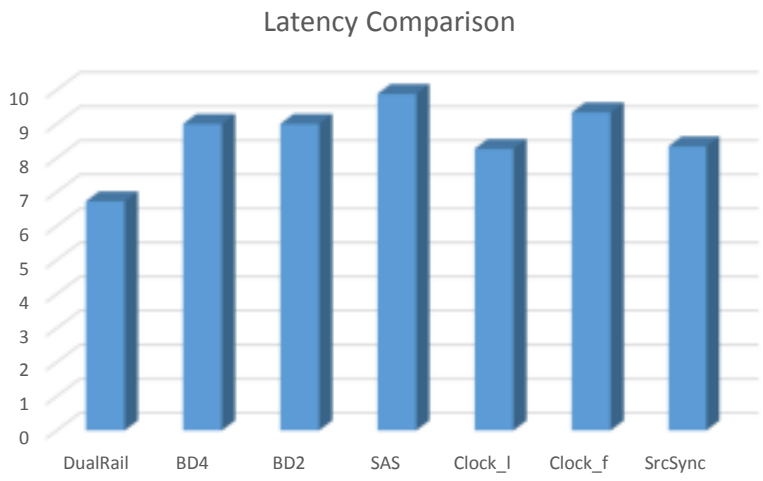

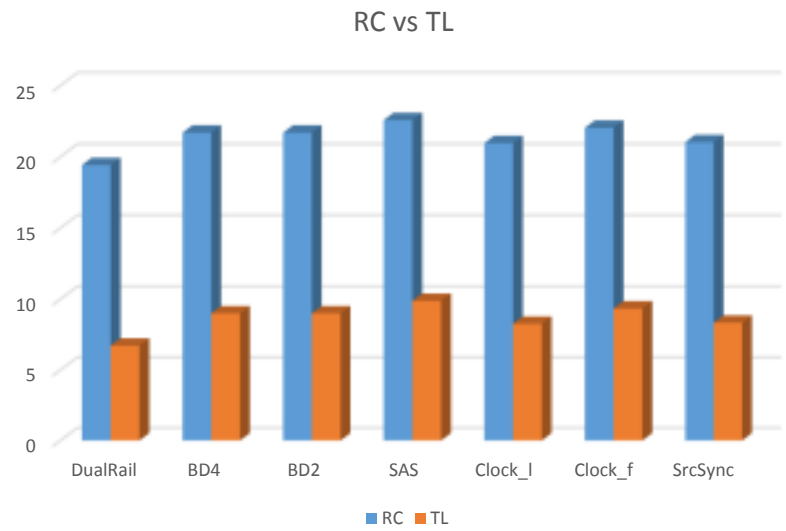

Comparisons - Energy

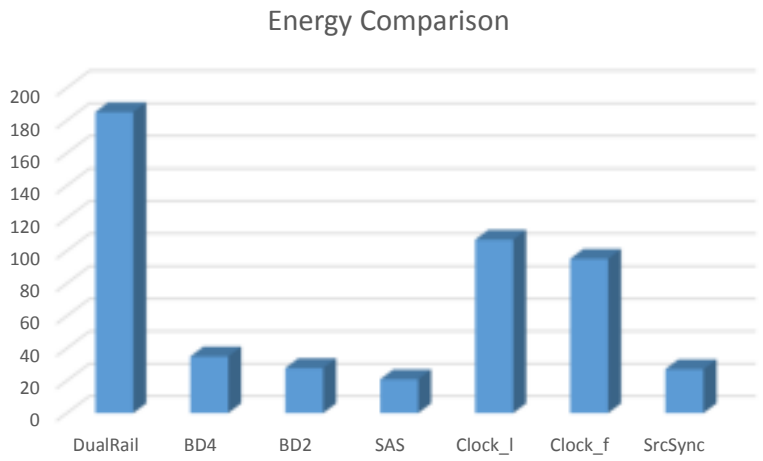

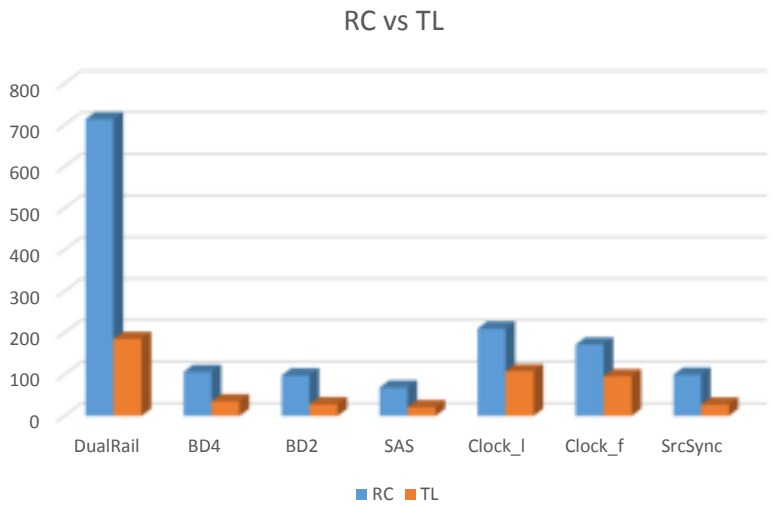

Key Observations

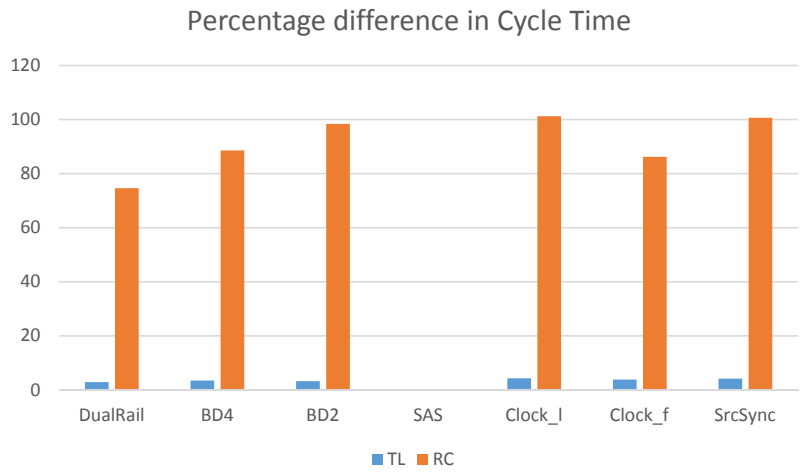

Effect of link length (7mm vs 3mm)

- Clocked protocols have better timing characteristics
- Clock distribution energy is a killer
- Single “cycle” communication due to discontinuity-free requirement of TLs
- SAS provides clocked-like timing without the energy overhead of clock distribution
- Longer distances more “manageable” using Transmission Lines
- SAS outperforms other protocols in almost all metrics
- No wavepipelining
- SAS robust to variation due to decoupled throughput and wire latency