



OPEN
Compute Project

BoW: A Die-to-Die Interface Solutions Specification Update

ODSA Project Workshop

June 10, 2019

Ramin Farjadrad



Mark Kuemerle

High-level Targets for a Inter-Die Connectivity IP in MCM Packages

- Performance Targets:
 - Throughput Efficiency: $\sim 1\text{Tbps/mm}$ (die edge)
 - Energy Efficiency : $1\text{pJ/bit} - 0.5\text{pJ/bit}$
- Small silicon area per IO port for dense integration
 - Goal: IP silicon area not to limit IO density
- Minimal analog/complex circuitry to offer easy/fast process porting
 - Limit the maximum baud rate of the interface : $\sim 10\text{G}-16\text{Gbaud}$
 - Common clock to replace CDR with DLL, Digital PLL to generate clock, etc
- Single supply IP supporting wide Vdd range: $0.70\text{V} - 0.9\text{V}$
 - To be compatible with most existing SoC/ASIC in popular/available process nodes
- Logical compatibility with AIB for ease of adoption

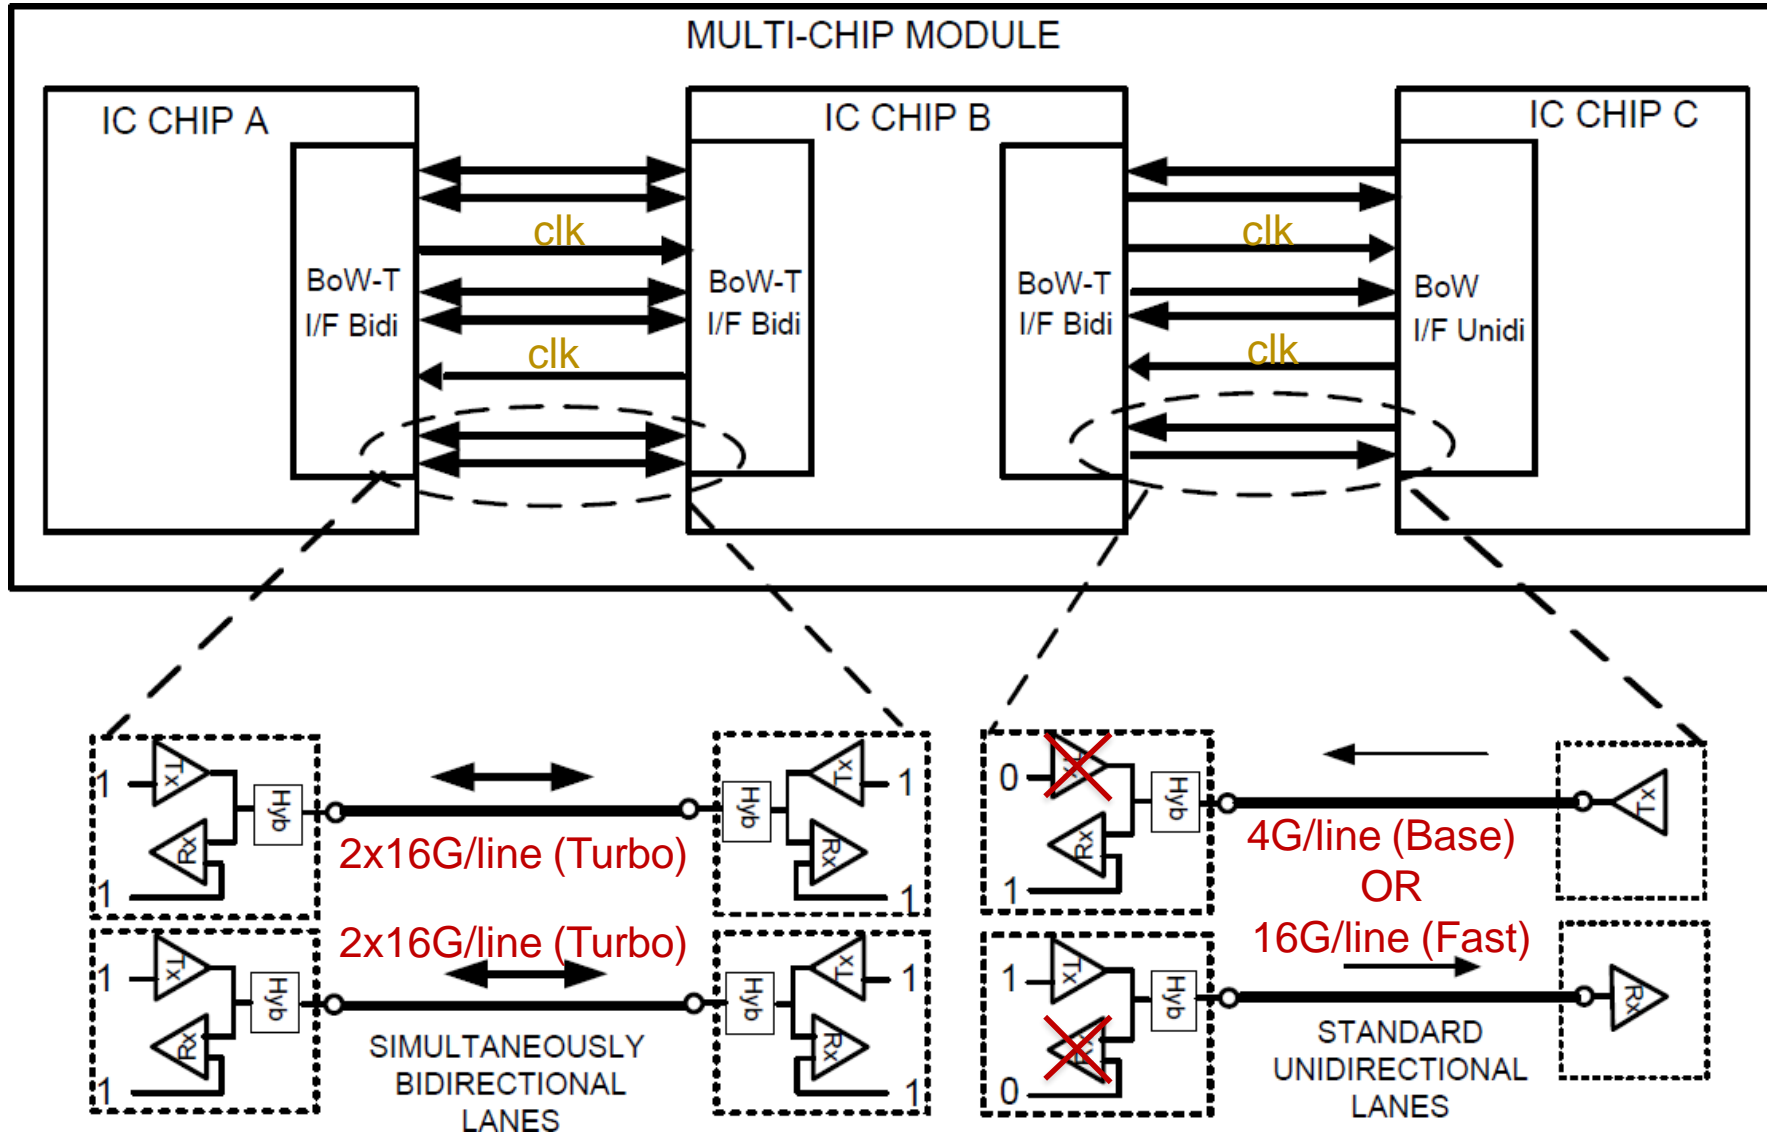
BoW (ODSA Proposal): Single-ended Signaling

- BOW Base
 - Unterminated lanes up to 5 Gbps/wire over 10mm
 - Source Synchronous with clock alignment
- BOW Fast
 - Terminated lanes up to 16Gbps/wire over 50mm
 - DDR Source Synchronous with clock alignment
- BOW Turbo
 - Simultaneous Bidirectional in both directions
 - Terminated up to 2x16Gbps/wire= 32Gbps/wire over 50mm
 - Source Synchronous with clock alignment




Note: Baud rate limited to 16Gbaud for simplicity of design and ease of port

Top Level View: Bow Backward Compatibility



- A Bow-Fast can be configurable to be backward compatible to Bow-Base
 - By disconnecting the line terminations
- A Bow-Turbo can be configurable to be backward compatible to Bow-Base/Fast
- Bow-Fast:
 - By disabling Tx or Rx per lane
- Bow-Base:
 - By disabling Tx or Rx per lane and disconnecting terminations

BoW Slice (Building Block) Bump Map

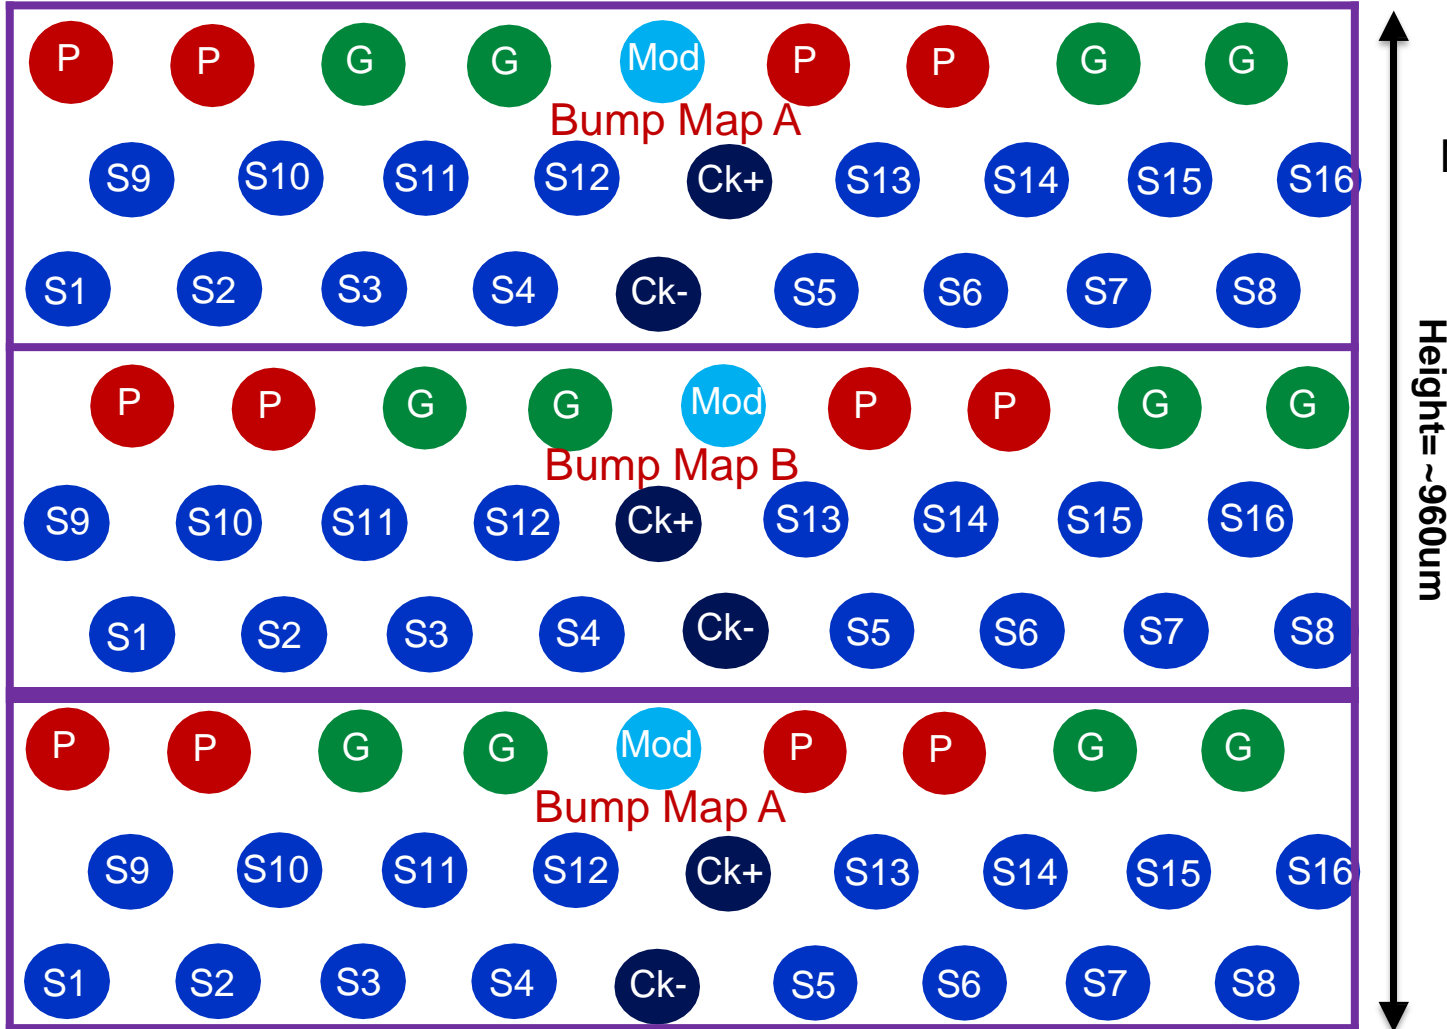


- BoW Slice with same circuitry/layout but different RDL comes in 2 bump maps to create efficiency in building larger BoW modules

- A common bump ordering to be used for **BoW Slice** Base/Fast/Turbo
- **BoW Slice**: 16 Data + 2 Clock ports
 - Configurable to be output clock or input clock pads when connected to non-Turbo Slices
- Vertical stacking of Bow Slices increases throughput per die edge
- **BoW Slice** Maximum Throughput
 - BoW Base :
 $16 \times 4\text{Gbps/pad} = 64\text{Gbps}$
 - BoW Fast:
 $16 \times 16\text{Gbps/pad} = 256\text{Gbps}$
 - BoW Turbo:
 $16 \times 2 \times 16\text{Gbps/pad} = 512\text{Gbps}$

BoW 3-Deep Module Bump Map

Chip edge = 1170um (Bump Pitch=130um)



- Stacking 3 (three) BoW Slices provides 48 data pads, with following throughputs:

Note: Throughputs are for aggregate of Rx + Tx

- BoW Base :

64x4Gbps/pad = 192Gbps

Throughput/mm = ~164Gbps/mm

- BoW Fast:

64x16Gbps/pad= 768Gbps

Throughput/mm = ~656Gbps/mm

- BoW Turbo:

64x2x16Gbps/pad=1540Gbps

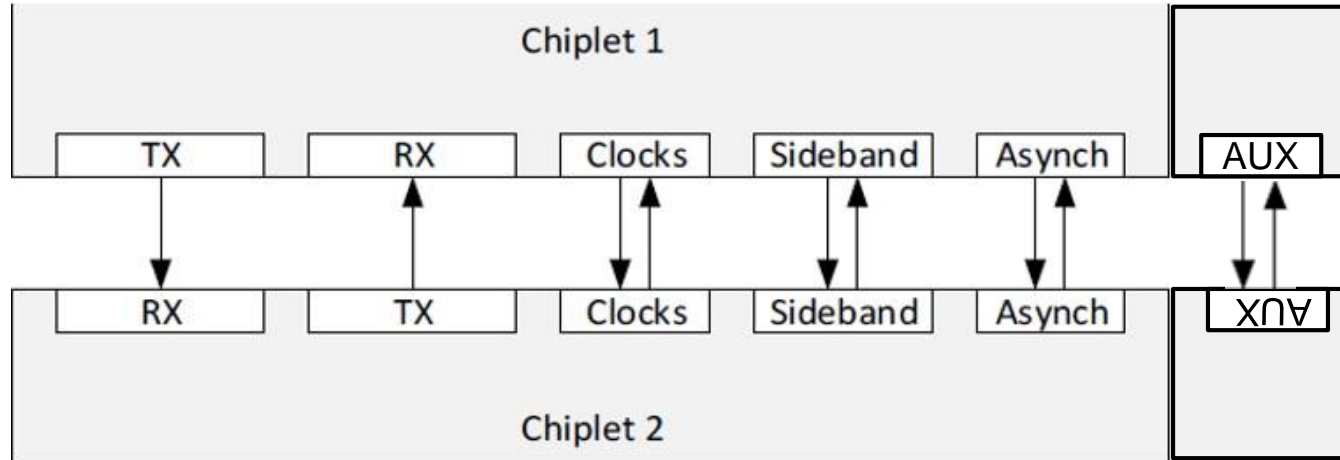
Throughput/mm = ~1.32Tbps/mm

- Needs a package with 3-2-3 stack-up

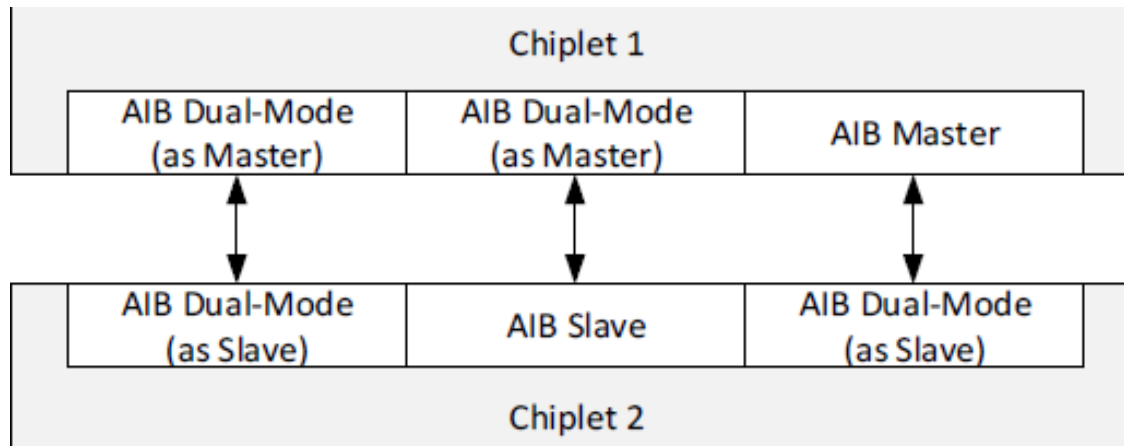
BoW Slice Specifications

Parameter	Parameter
Single Supply Voltage	0.70V-0.90V (+/-5%)
Baud rate/bump	Base: 1-4Gbaud, Fast: 1-16Gbaud, Turbo: 1-16Gbaud
Max Throughput/mm (Combined Tx+Rx)	Base: 55Gbps, Fast: 220Gbps, Turbo: 438Gbps ← 1x Stacked Base: 110Gbps, Fast: 438Gbps, Turbo: 875Gbps ← 2x Stacked Base: 164Gbps, Fast: 656Gbps, Turbo: 1320Gbps ← 3x Stacked
BoW Slice Dimensions (Bump Pitch=130um)	Chip Edge: 1170um, Height:320um
Energy Efficiency (14nm)	< 0.7pJ/bit
Trace length	Base: 10mm, Fast: 50mm, Turbo: 50mm
Latency	< 30/Gbaud (<3ns @10Gbaud)
Slice Bump Allocation (excludes global AUX)	16 Data, 2 Clock, 1 Mode, 4 Ground, 4 Power
BER Target	<1E-15 (No ECC)
ESD / CDM protection	400V/100V

BoW Compatibility with AIB





AIB Signal Types




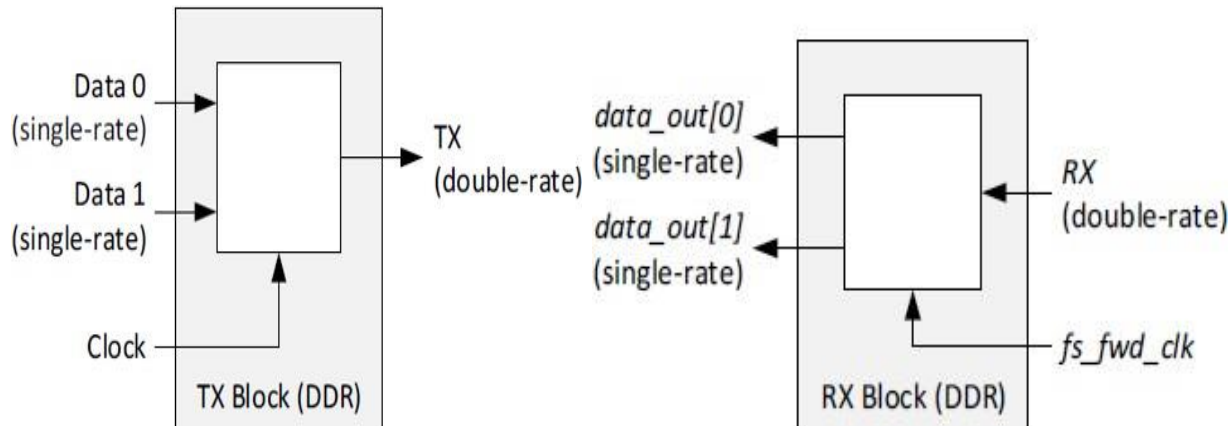
Dual Mode Interfaces

- From logical viewpoint, BoW to use similar Interface Signals as AIB, with following differences:
 - Master Interface to also be the timing Master
 - Sideband Controls limited to 32 registers for both Master and Slave
- AUX signals are not duplicated like AIB and are independently hardwired
- BoW to provide Dual mode Master and Slave option
- 16 bit bus vs. 20+ bit bus

SDR & DDR Clocking




SDR




DDR

Receive SDR and DDR I/O Mapping

DDR Mode




SDR Mode




- Clock signal `hs_clk` is fed to BoW IO block from a clock generator (e.g. PLL)
- In Bow Basic mode, `fwd_clk`, `hs_clk`, and `mac_clk` all run at the same speed
- In BoW Fast and Turbo, `fwd_clk` and `hs_clk` run 4x faster than `mac_clk`
- Mux/Demux ratios will be different between BoW Basic (2:1/1:2) and Fast/Turbo (8:1/1:8)

Transmit SDR and DDR I/O Mapping

DDR Mode



SDR Mode



- Clock signal mac_clk is always sent from BoW IO block to the MAC block, regardless of the block being a receive or transmit IO
- Clock signal hs_clk is fed to BoW IO block from a clock generator (e.g. PLL)

BoW Sideband Control Signals

Signal name	Signal function	Bits	Signal origin (far-side or MAC)
Calibration (Section 3.2.3)			
<i>ms_osc_transfer_en</i> <i>sl_osc_transfer_en</i>	Oscillator calibration complete	1	FS
<i>ms_tx_dcc_cal_done</i> <i>sl_tx_dcc_cal_done</i>	TX DCC calibration complete	1	FS
<i>ms_rx_transfer_en</i> <i>sl_rx_transfer_en</i>	RX calibration complete	1	FS
<i>ms_rx_dll_dcc_lock_req</i> <i>sl_rx_dll_dcc_lock_req</i>	Start RX calibration	1	MAC
<i>ms_rx_dll_lock</i> <i>sl_rx_dll_lock</i>	RX DLL locked	1	FS
<i>ms_tx_transfer_en</i> <i>sl_tx_transfer_en</i>	TX calibration complete	1	FS
<i>ms_tx_dll_dcc_lock_req</i> <i>sl_tx_dll_dcc_lock_req</i>	Start TX calibration	1	MAC
User Defined			
External Control	Defined by protocol or application	MS: 8 SL: 8	FS or MAC
Other			
Reserved		MS/SL:9	NA

- BoW uses the same 7 fixed Sideband Controls as AIB, 16 external controls for both Master & Slave, and 9 Reserved controls
- Limit sideband clock to 200M-400MHz
 - Fewer controls bits relaxes clock frequency
- The functions of the fixed sideband controls are the same as AIB with same source:
 - Far-side (link-partner)
 - MAC (local)

Test & Calibration Options

- 1149.1 Legacy/1149.6 High-speed (JTAG scan)
- IEEE 1500 (HBM type systems)
- At-Speed Self Test
 - Serial At-Speed PRBS Self Loopback: Tx(Port-N_Chiplet-A) to Rx(Port-N_Chiplet-A)
 - High-Coverage Wafer Test Screen (pre-package) loopback traces on ATE load board
- At-Speed System Test/Compliance
 - External PRBS Loopback: Tx(Port-N_Chiplet-A) to Rx(Port-M_Chiplet-B)
 - Eye Monitor for At-Speed Test/Compliance: Tx(Port-N_Chiplet-A) to Rx(Port-M_Chiplet-B)
 - Per data bump: Measure Errors Rate for each phase and voltage threshold
- Calibrations
 - PLL/DLL Lock
 - Rx Phase Lock
 - DCC Calibration (optional)

Summary

- Concept proven in 14nm Silicon (Hybrid easy to port to other nodes)
- Over 1Tbps/mm chip edge over 50mm organic substrate (pitch= 130um)
 - Full-duplex Throughput/bump up to 32Gbps (2x16Gbps)
- Small area per port of $<0.02\text{mm}^2$
- Less than 0.7pJ/bit in 14nm ($<0.5\text{pJ/bit}$ estimate in 7nm)
- Single power supply 0.7V-0.9V: Compatible with synthesized logic circuits
- Easy and quick to port into other process nodes (16G NRZ vs 112G PAM4)
- Backward compatibility
 - A Chiplet with Bow Turbo interoperates with other BoW interfaces
- High Wafer-level test coverage per Chiplet to improve final product yield


Target Specification Milestone

- We Are Targeting:
 - Spec Draft V0.5 by Hot Interconnect Workshop in August
 - Spec Draft V0.9 by next ODSA Workshop in September
 - Spec Draft V1.0 by ODSA Workshop in November

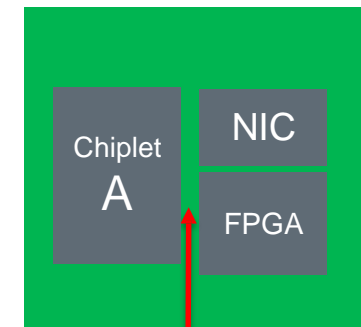
Why BoW Specification Makes Sense

- There are multiple versions of die-die interfaces being built by multiple companies today
 - Open standard makes sense to converge design points
- A common and simple parallel interface at relatively low data rate enables chips in older nodes to be integrated into the system
 - Example: Critical for RF component integration
- Master/Slave mode simplifies clock generation requirements
 - Also helpful for legacy technology implementations
- Expandability of the same pins in higher data rates with BoW Fast and Turbo allows for chiplets to meet multiple applications
 - Example: Streaming data with several RF chips

Same die switching accelerator data with NIC, FPGA



Hundreds of Gbps




Thousands of Gbps

How You Can Help


Seeking volunteer(s) for

- Foundries to support implementation of BoW PHY IP in their high-end process nodes (esp Finfet)
- Development of the BoW PHY IP in multiple process nodes (28nm to 7nm and even 5nm)
 - Availability of the BoW in wide range of process nodes will encourage larger and faster adoption
- Development of PHY Interface to different Standards and Protocols
 - BoW PHY needs a separate dedicated interface layer developed for each transmission standard having different coding and partitioning
 - Example: PCIe/CXL, Ethernet, etc






BACK UP

AQlink Demo Silicon in GF




LR: M1 ↔ S2	MR: M2 ↔ S1	SR: M3 ↔ S0
2mm	10mm	25mm



traces on test chip

Consume. Collaborate. Contribute.

AQlink Internal Eye for Bidirectional 28Gbps ($>1E^{11}$ bits per point)



Consume. Collaborate. Contribute.