Compute Project

BoW: Basic, Fast, Turbo Die-to-Die Open Interface Solutions

ODSA Project Workshop

March 28, 2019

Mark Kuemerle

Ramin Farjad

Consume. Collaborate. Contribute.

Interface standardization

- Chiplet based systems will likely have a variety of different interfaces moving forward
  - Fast Serial interfaces with some compatibility to existing standards (eg 112GXSR)
  - Novel USR interfaces driven by an ecosystem (eg. Kandou/Marvell)
  - **Simple interfaces for moving reasonable bandwidth between die from less cutting edge nodes**

Does it make sense?

- Does it make sense to have another parallel standard?
  - AIB is a good starting point, currently has datarate and footprint definitions that limit use on a laminate
  - Making the spec as open as possible can hopefully speed up useability

- How do we decide if this effort is providing a value to the community?
  - Inertia from participants, adoption in products
    - Standardization coalesces
    - Adoption in custom/'contained' interfaces
    - Moving to interoperable interfaces between various die

The basic idea

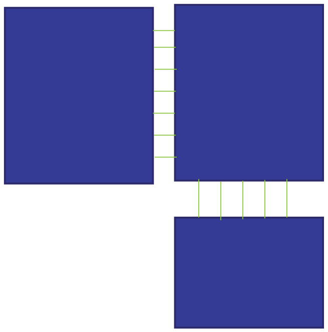

- Blast from the past – use simple CMOS IO to communicate
- We stopped using these when SERDES simplified board routing (less traces=less $$)
- Keeping everything on a laminate keeps things simple
  - More routing traces than a board
  - Less ESD requirements
- Make an interface that works on a cheap laminate or a fancy silicon based interconnect
  - Enable compatibility where we can

BoW Intro:

- High level proposed solution:
  - Simple source synchronous DDR interface
    - Clock adjust needed for DDR
  - Non Terminated (Termination may be ok, just adds power at lower rates)
  - 1-4 Gbps (or more?)
  - Low overhead IO cell (limited ESD)

BoW Extensions:

- Data Integrity / Power functions – need to discuss
  - ECC – How critical is BER? How much does interface spend vs.
  - DBI – is power worth the pins?

- Extension Mode 2: Bow-Fast
  - Higher speed uses optional termination (~2-3x rate), minor impact on power

- Extension Mode 3: Bow-Turbo
  - Add potential bidirectional feature – implemented in the IO cell
    - Hybrid circuit – can be implemented in multiple ways (examples from BaseT)

Proposal for dart throwing

Function	Pins	Notes
RX Data	32	
Data Clock	2	Differential
Parity	1	

Function	Pins	Notes
TX Data	32	
Data Clock	2	Differential
Parity	1	

Minimum streaming IF

Other Function	Notes
Data Bus Inversion	~19% power savings
Parity	Not needed if ECC
FIFO Reset	Helpful for control/init
ECC	Error detect/correct
Sideband	Various calibration/etc

Other Function	Notes
Data Bus Inversion	~19% power savings
Parity	Not needed if ECC
FIFO Reset	Helpful for control/init
ECC	Error detect/correct
Sideband	Various calibration/etc

Datarate discussion

- Current proposal: Per databit wire configurable from 1-4Gbps
- What datarates make the most sense for the standard?
  - 1-2Gbps (AHB) limits usefulness on laminate (too many wires to move reasonable bandwidth
  - Configurable datarate including 1-2 Gbps could enable AHB compatibility
  - >5Gbps may require termination
- Having a settable datarate with simple divides maximizes compatibility
  - For 4, 2, 1 Gbps -> simple dividers

Voltage discussion

- Simple, single voltage based approach – challenge is what voltage to choose?

- 1.2 is HBM legacy, costs power
- 0.9V is often available, decent power, interoperable with AIB*
- 0.75 / 0.8 often available as well, popular chiplevel VDD values
- Lower voltages will have lower power, too low will be a challenge for IO design

Power (energy)

- Simple energy calculation
  - Pathological worst case w/DBI
    - 0.4
  - No on chip routing from IO to bump assumed
    - Add 0.1 – 0.25 pJ depending on length
  - Termination +~0.2 pJ low speed, less at high speed

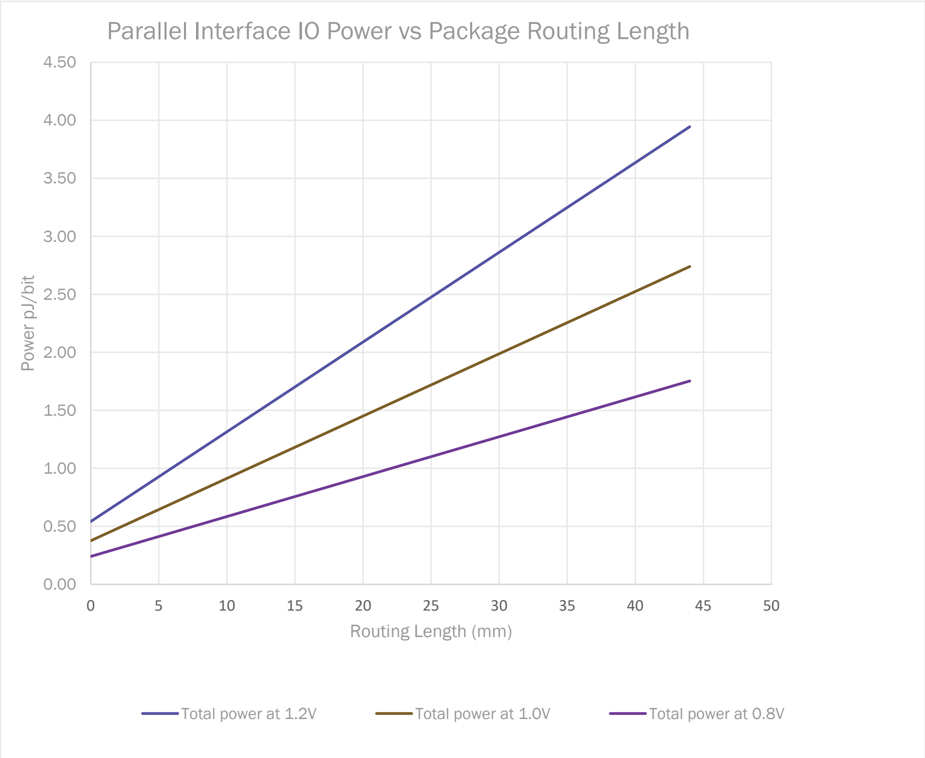

OPEN Compute Project

Operation Modes BoW on Organic Package Substrate

- BOW Basic
  - Unterminated lanes → up to 5 Gbps/wire
  - Source Synchronous with clock alignment

- BOW Fast
  - Terminated lanes → up to 12Gbps/wire
  - Source Synchronous with clock alignment

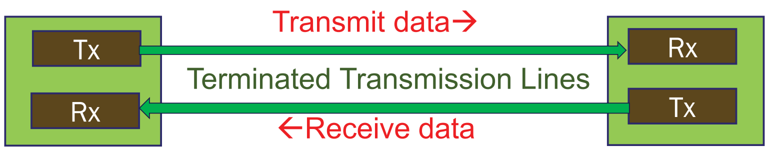

- BOW Turbo
  - Simultaneous Bidirectional → both directions
  - Terminated lanes → up to 2x12Gbps/wire
  - Source Synchronous with clock alignment

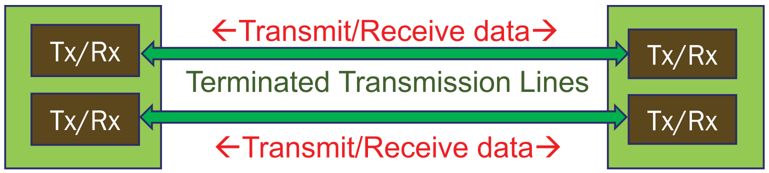

Note: Bidirectional signaling been around for decades, in all phone lines, and at multi-Gig in BASE-T PHYs since late 1990s. The cancellation requirement for a link with small loss is fairly relaxed

BoW-Turbo IO Block Diagram

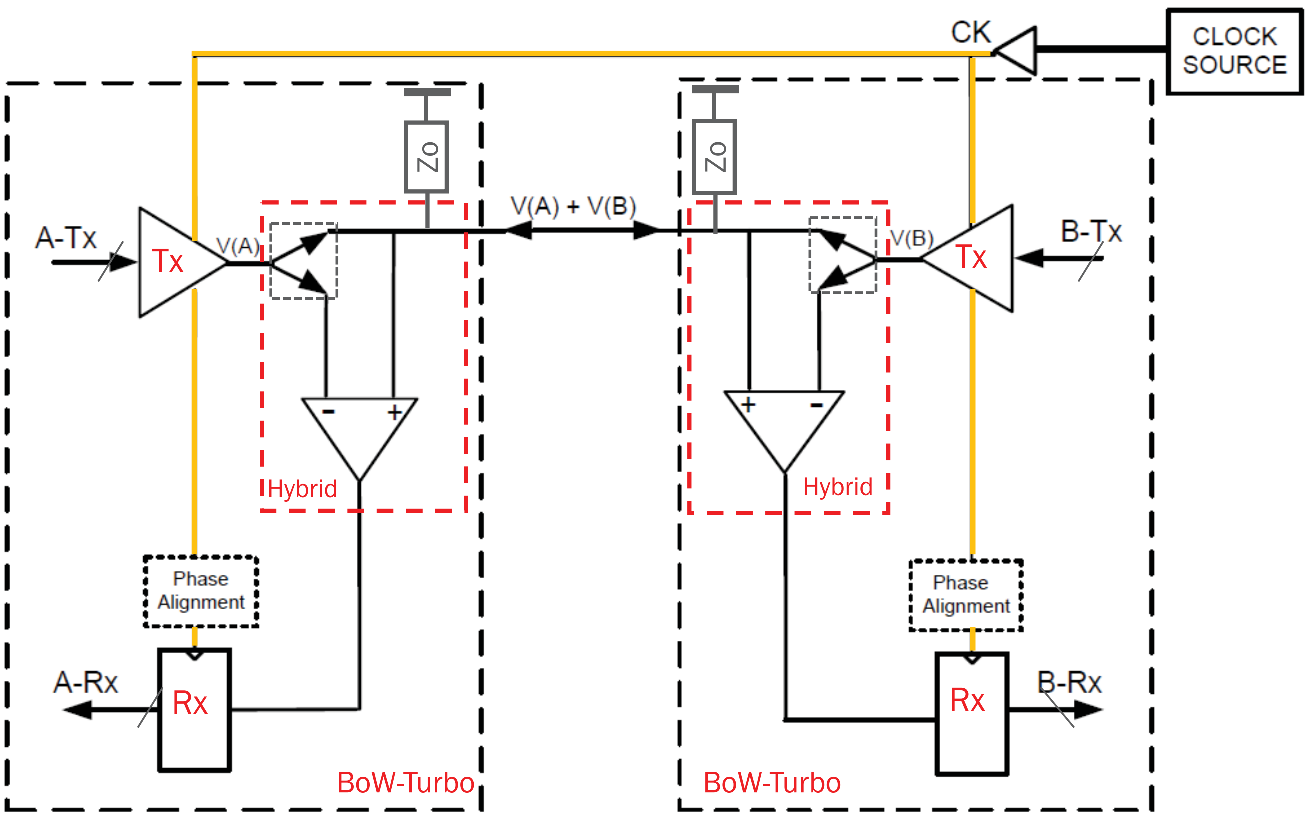

- A Hybrid block, placed between the pad and basic BoW Tx/Rx ports, creates a Bow Turbo port
  - Passes local Tx output signal to the IO pad
  - Subtracts local Tx output signal from the combined signal on IO pas and passes the resulting signal to Rx input
  - Once implemented, Easy to port
- BoW-Turbo is backward compatible with Bow-Basic & BoW-Fast
  - Interface can easily be programed to act as a transmit only or receive only port
- Total area for BoW-Turbo transceiver is <0.018mm^2.
  - Can be integrated under single pad area (pad pitch=130um)
- Bow-Turbo concept proven in Silicon
  - AQlink SerDes in GF 14nm operated in Single-ended mode

AQlink Demo Silicon in GF 14nm

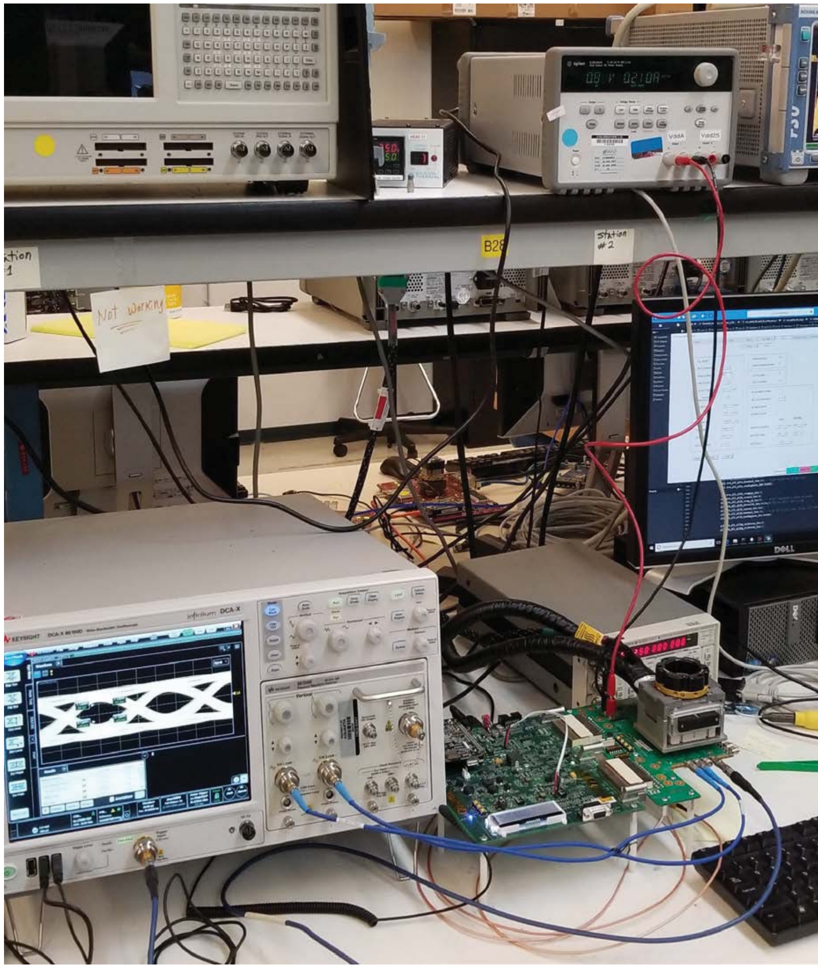

LR: M1 ⇔ S2	MR: M2 ⇔ S1	SR: M3 ⇔ S0
2mm	10mm	25mm

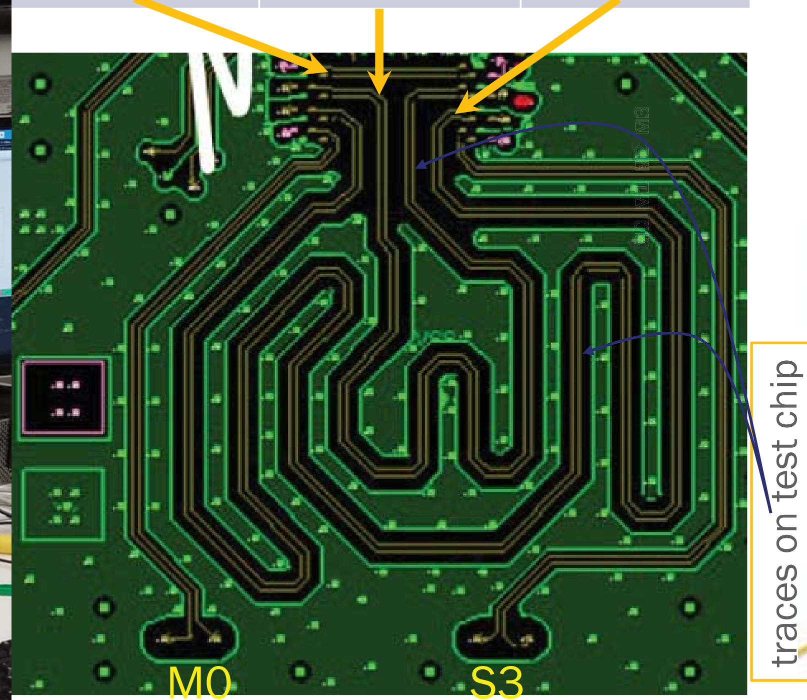

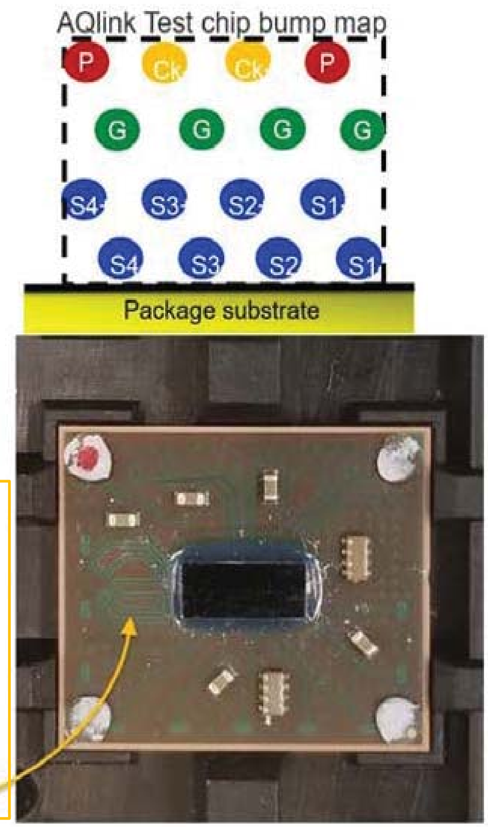

AQlink External Ports

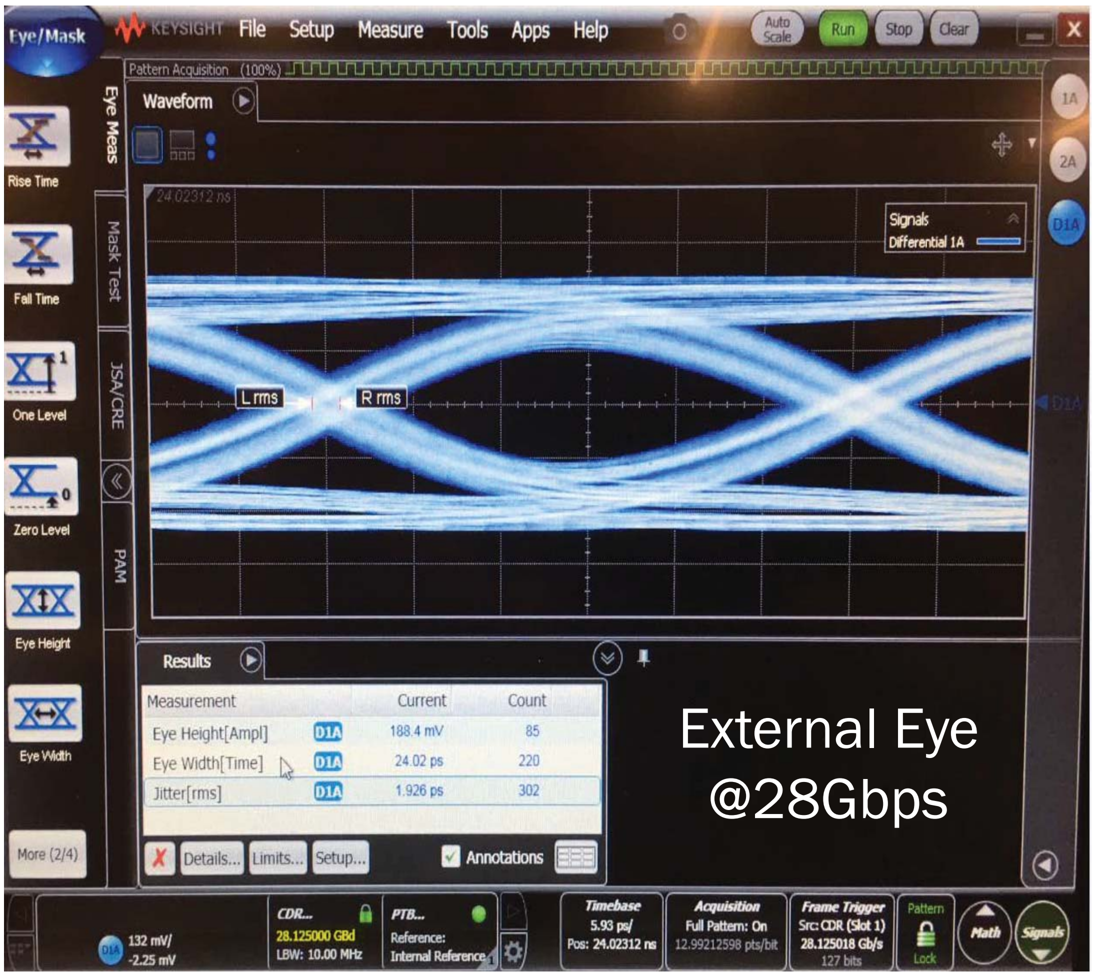

Consume. Collaborate. Contribute.

- One port per AQlink Quad core is an external port. Each external port trace:
  - 10mm package
  - 30mm board (1/2 package loss)
  - →Effective 10+15=~25mm package
  - →Loopback trace=2x25=50mm
- A differential external loop back operates error free up to ~25Gbps per direction (BER~1E-15)
- If loopback connected in single-ended fashion, the link operates error free up to ~15Gbps/direction (BER~1E-15)
- To provide further margin, simplify the phase alignment & duty cycle corrector, maximum Bow-Turbo speed limited to 12Gbps/direction

AQlink Internal Eye for Bidirectional 28Gbps (>1E11 bits per point)

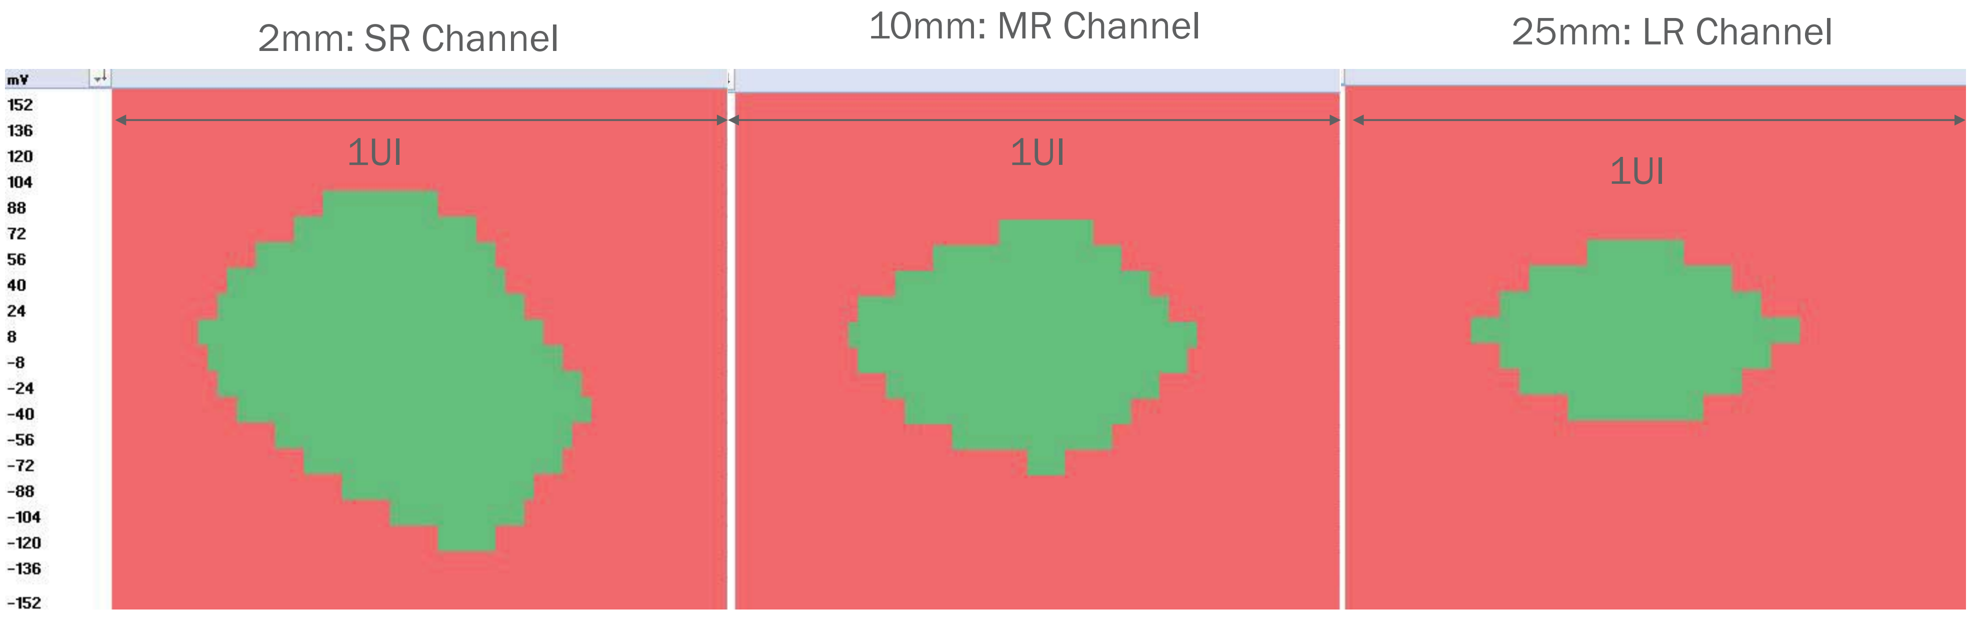

Top Level View: Bow-Turbo Backward Compatible to Fast/Basic

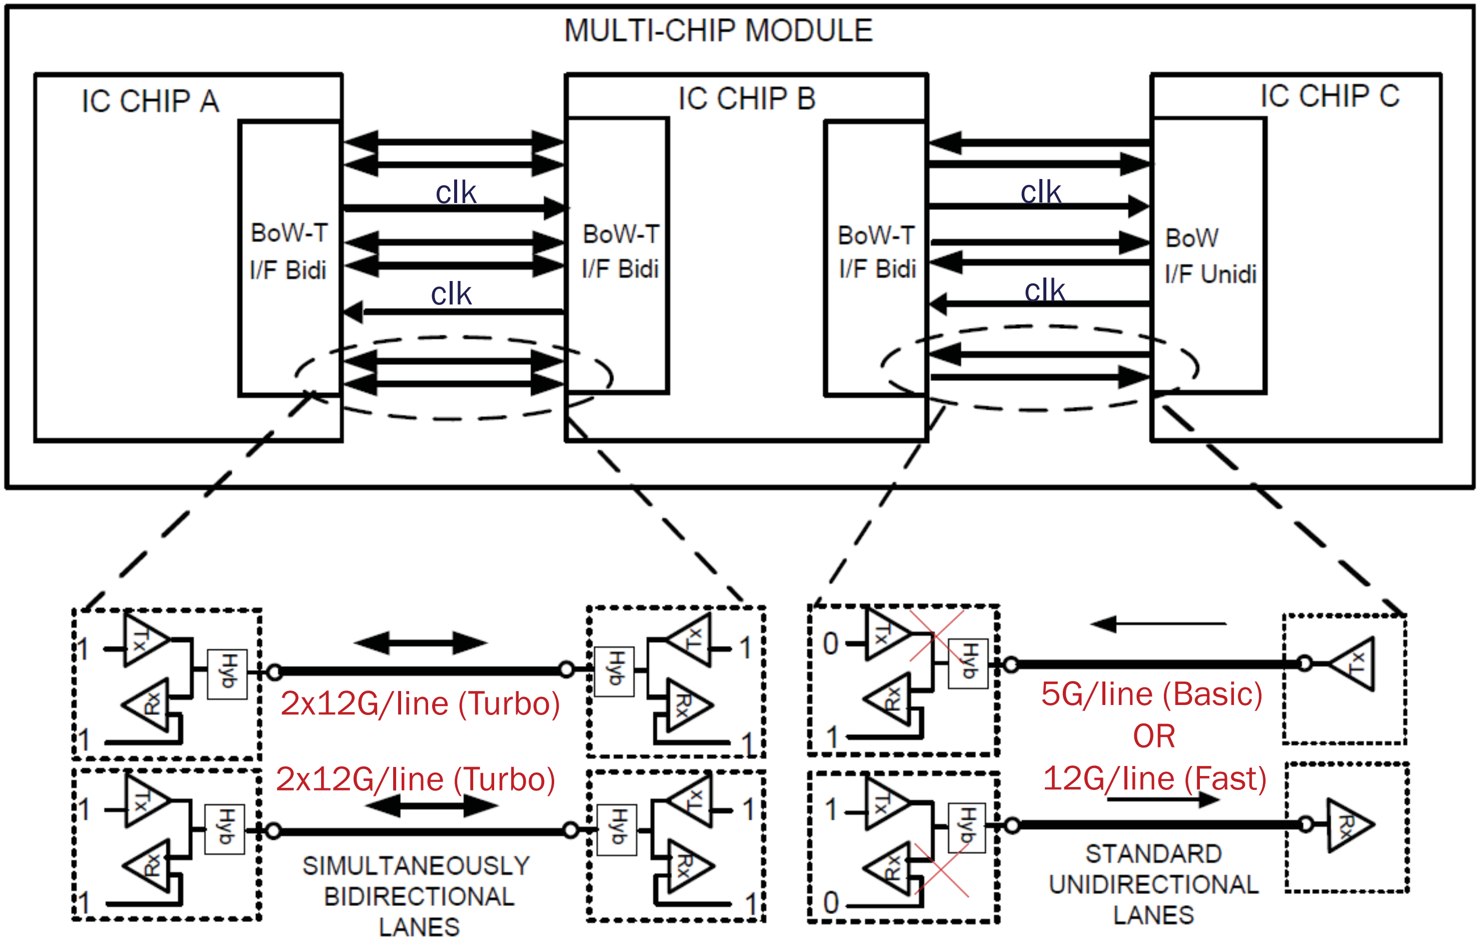

- A Bow-Turbo Core is configurable to be backward compatible to
  - BoW-Fast
    - by disabling Tx or Rx per lane
  - BoW-Basic
    - by disabling Tx or Rx per lane
    - disconnecting the line terminations

Sample BoW Staggered Bump Map – Terabit Core

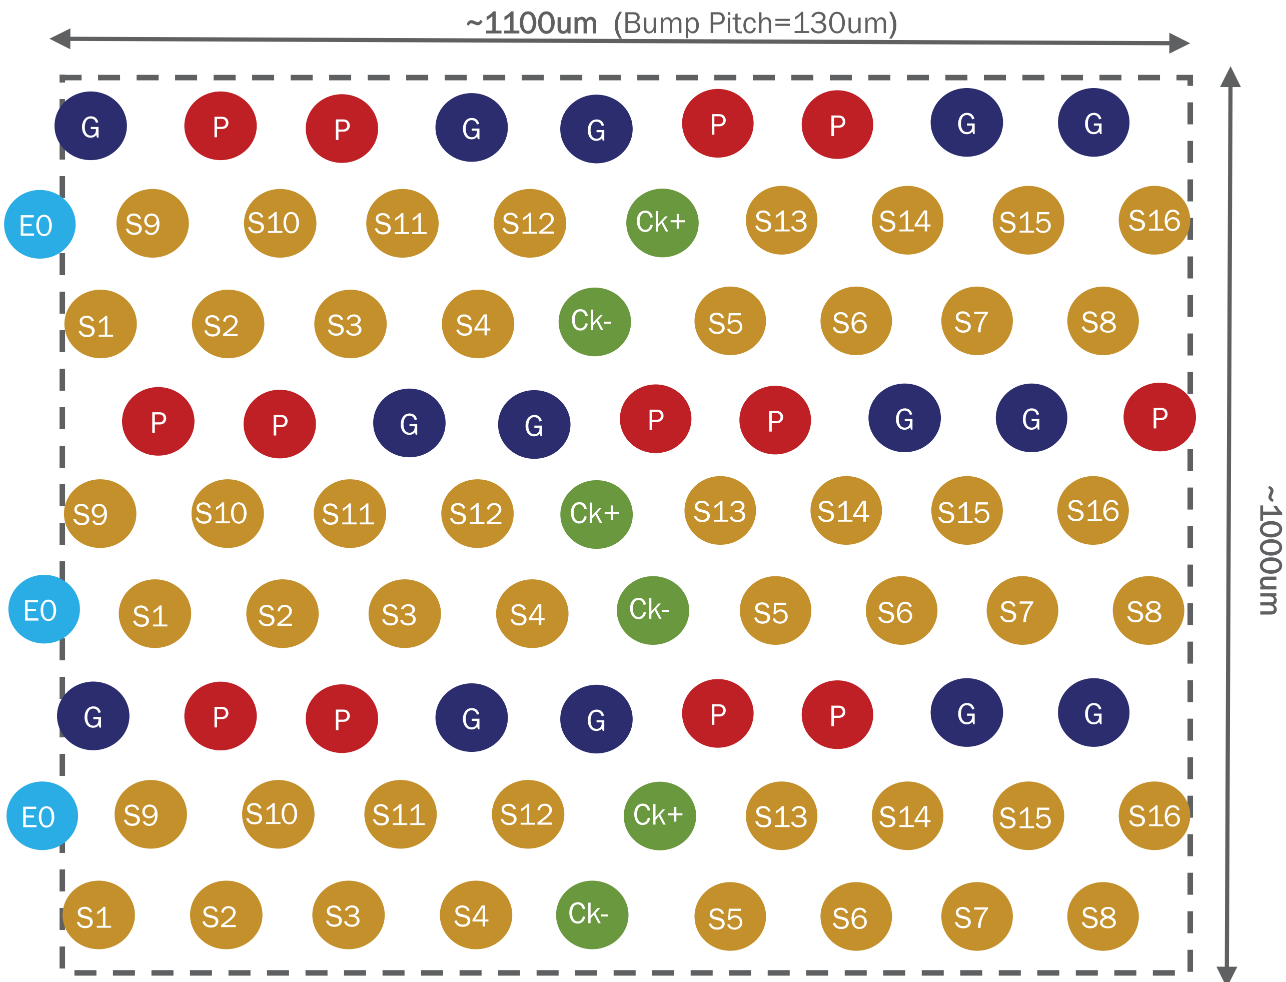

- Terabit Module has 48 data pads:
  - 48x2x11.5Gbps/pad=1Tbps/mm
- A common bump map can be used for Basic/Fast/Turbo
  - Important for backward compatibility
- Two clock ports per 16 data ports
  - Configurable to be clock output or clock input pads when connected to non-Turbo interface (Tx or Rx only)
- Optional ECC
  - Option column to add one ECC bit per 16 data group at the edge of the Terabit Core bump map to provide error correction capability if BER<1E-20

Power Efficiency of BoW Fast/Turbo

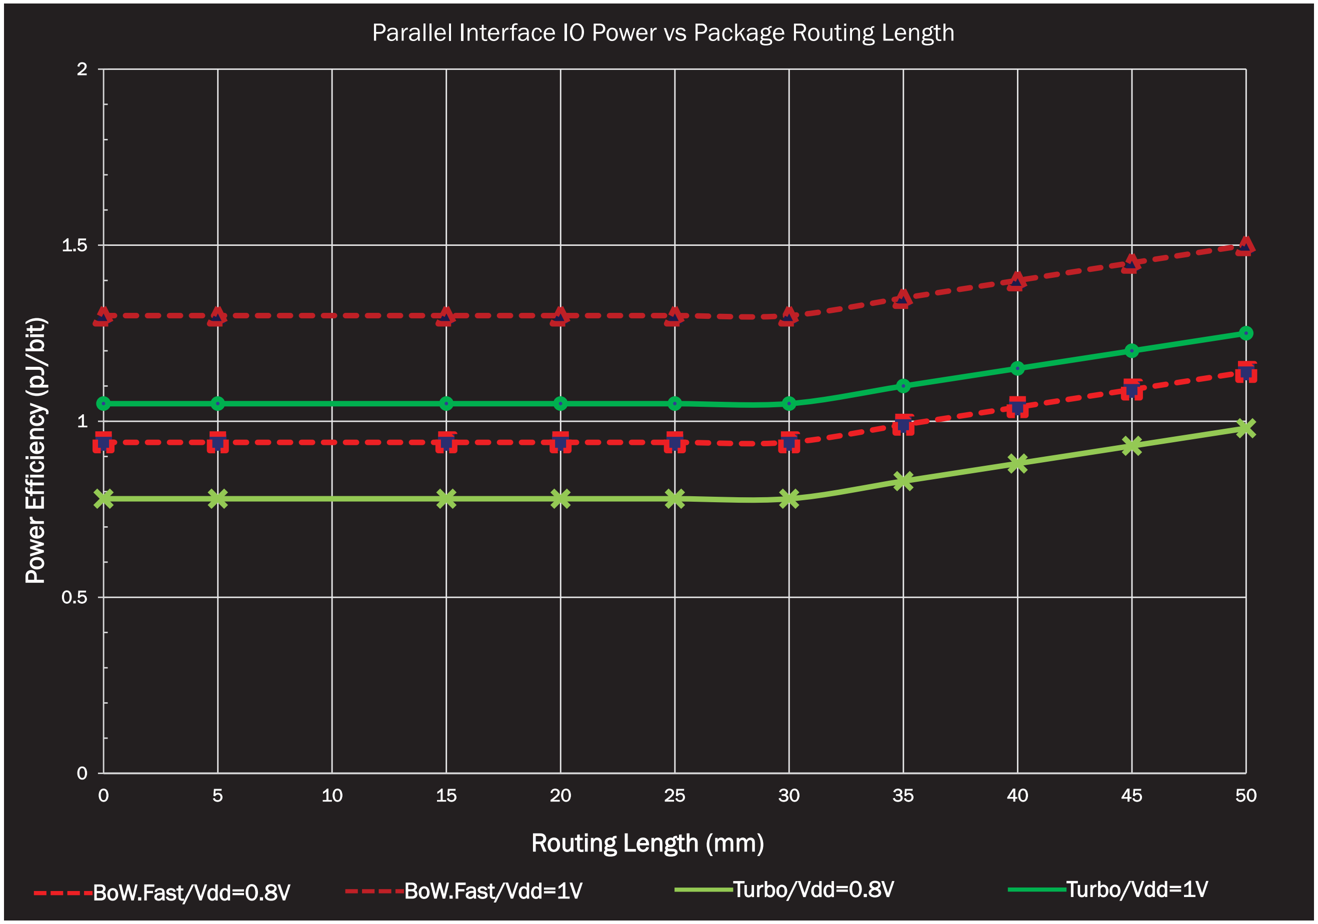

- Data activity → PRBS
- Power efficiency for data from the edge of the interface core on side to edge of interface core on the other side
- Based on 14nm 28Gbaud silicon measurement (AQlink), but more room for power reduction:
  - Port from 14nm to 7nm
  - Baud reduction from 28Gbuad to 12Gbuad means less circuitry

BoW-Turbo Interface Performance Specifications

Parameter	Parameter
Single Supply Voltage	0.75V-1.2V (+/-5%)
Throughput/Trace (Max)	24Gbps (2x12Gbps)
Power Efficiency	0.74pJ/bit (0.8V/30mm/14nm)
Package Trace length (Max)	50mm (Package Substrate: GZ41)
Latency	<2ns
Pad Pitch	130um
Terabit IP Core Dimension	Chip Edge: 1100um. Height:1000um
Power/Area for 1Tbps Throughput	740mW/1.1mm^2
BER	<1E-15 (No ECC) / <1E-20 (with ECC)
ESD / CDM protection	400V/100V
Silicon Proven	GF 14nm

BoW Turbo Spec Summary

- Over 1Tbps/mm chip edge over organic substrate & 130um pad pitch
  - Throughput per port of 2x10.5Gbps (up to: 2x12Gbps)
  - Small per port area of <0.018mm^2
- Less than 1pJ/bit in 14nm at Vdd=0.8V and trace=50mm
- Single power supply that is compatible with synthesized logic core
- Concept proven in 14nm Silicon (Hybrid easy to port to other nodes)
- Easy and quick to port into other process nodes (Just a 10G USR)
- Backward compatible with BoW Basic/Fast die-die interfaces
  - Can use matching bump map as Bow Basic/Fast
  - A Chiplet with Bow Turbo interoperates with all Chiplets using other BoW interfaces

Test Options

- Need discussion to find suitable test solutions
  - Calibration?
  - 1149.1 legacy (JTAG scan)
  - IEEE 1500 (HBM type systems)
  - Could define at speed / functional interop test if needed

Call for Volunteers

- We need your help to define a generally useful and interoperable interface

- What do we need?
  - Good ideas on what applications the BoW makes sense for
  - SI/PI input for various options
  - Defining test for interface

Thank You

OPEN Compute Project

Consume. Collaborate. Contribute.