PHY Layer

ODSA Project Workshop

March 28, 2019

Consume. Collaborate. Contribute.

Contributors

- Bapi Vinnakota – Netronome
- Greg Taylor – zGlue
- Halil Cirit – Facebook
- Ramin Farjad – Aquantia
- Farzin Firoozmand – Alpha Wave
- Brian Holden – Kandou
- David Kehlet – Intel
- Marc Kuemerle – Global Foundries
- Ralph Page – Samtec
- Tony Pialis – Alpha Wave
- Marc Verdiell – Samtec

Goals

- Propose a standard physical layer for delivering high data rate connections between chiplet products
  - Enabling chiplet products from independent sources and heterogeneous technologies to operate together in a single SiP
  - With low overhead from inter-chiplet bandwidth, latency, and power concerns

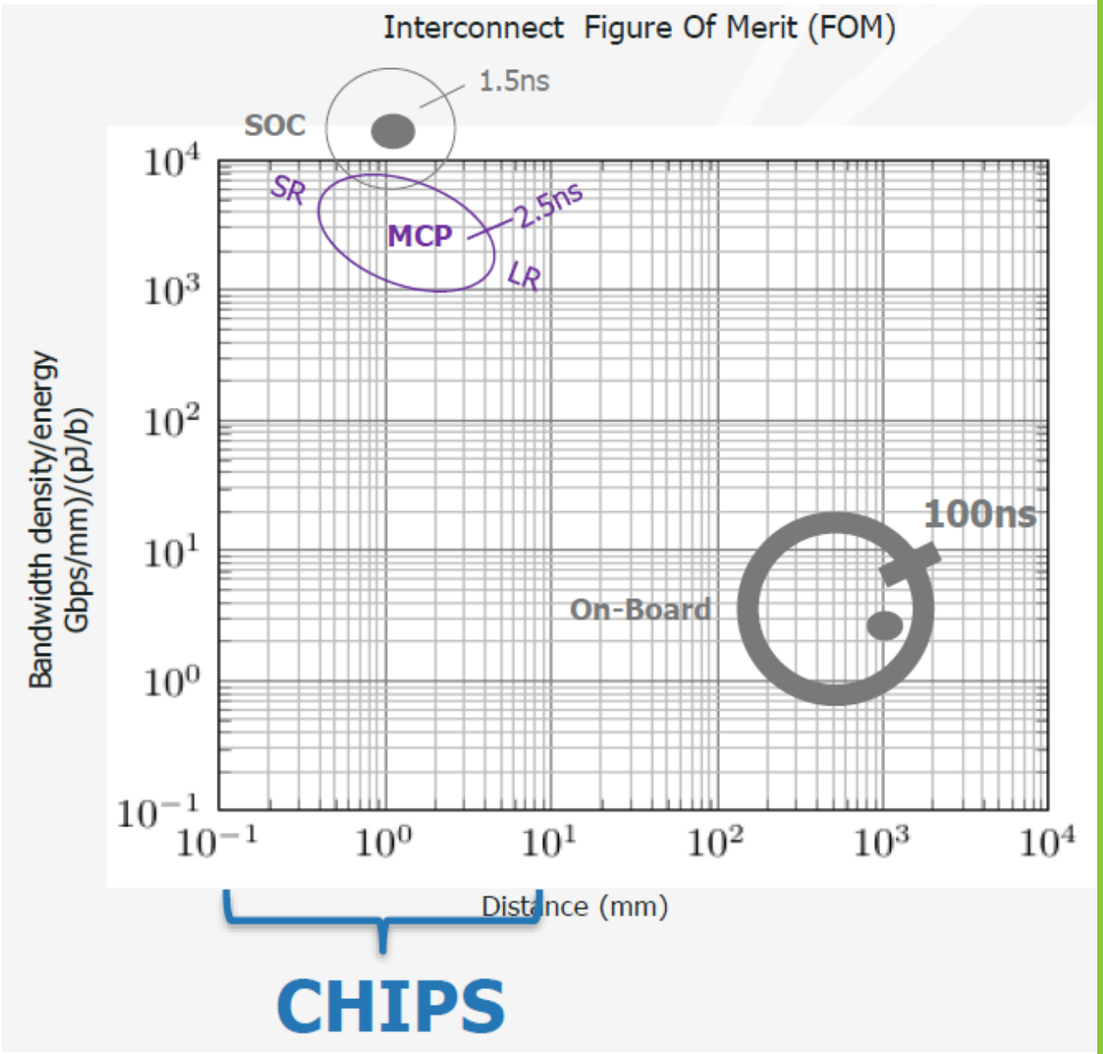

Process

- Generate a list of quantitative and qualitative interface concerns
- Generate a list of potential interfaces
- Gather technical data addressing the concerns
  - Also gather opinions on the qualitative concerns and the relative importance of the different parameters
  - Recognizing that different markets may have different priorities
- Develop proposals

Status

- Seven possible interfaces have been suggested
  - A mix of parallel and SERDES based approaches
  - A mix of Si interposer and organic packaging options
  - Both open and licensed
  - We have gathered quantitative data across the range
  - There is no obvious winner
- Some market divisions have been identified
- Now evaluating more qualitative concerns

Global Design Constraints

- Assume each interface supports multiple discrete data rates and widths
- SAM - $500M for any chiplet-based product – the intent is to build custom parts for targeted applications
- After manufacturing all interfaces need to support some element of end2end link self-test for both functionality and performance
  - Self-calibration only after chiplet assembly, can not require extra calibration after chiplets are assembled
  - Link should support stress tests and time/voltage margining
    - How do we learn from PCIe?
- Assume the test/control/management interfaces are likely to be BOW
  - Default to the JTAG bus as the OAM interface for any data interface?
  - Do we mandate a separate functional test interface – a particular form of JTAG?

Assumptions

- Open standards (XSR) will have multiple IP suppliers
- Interposer is a huge barrier for small-budget
- BOW Basic wins for very-low throughput markets

Metrics for Comparison

- **Piece Cost per Unit** - Manufacturing cost/complexity
  - Pad limited area – using area as a proxy for cost, assuming that the interface won't be used in process nodes where it is not pad limited
  - Substrate type – silicon or organic
- **Operating Cost** - Power-performance at rated throughput
  - Figure of Merit - DARPA definition – (Tb/mm)/(pJ/bit) – (die_edge_signal_density)/power_efficiency)
  - Technology node at which FOM is being computed
  - (In + Out) Bandwidth/mm as a proxy for max bandwidth in the package

Metrics for Comparison (cont)

- **Chiplet/Product Design NRE/Schedule RIsk** - Use cost/complexity
  - Routing freedom – as a proxy for design complexity (Number of wires impact (e.g. turning corners), length restrictions, pad placement restrictions, die placement restrictions
  - Ability to go in/out of low power states – to support chip power states
  - Proxy for number of process nodes diversity - TSMC node at which the circuit and pad area are balanced given a 130µm pad pitch
  - IP integration complexity – likely to be a wash – greater complexity in chiplet integration may make product integration easier
  - Product test and assembly – impact of interface/substrate on ability to test KGD
- **Interface Technology NRE/Schedule Risk** - Interface Design cost/complexity
  - Licensing fee to use the interface specification
  - Multi-sourcing – assuming multiple competition lowers costs
  - Interface test and assembly – estimated complexity to self-test functionality and performance of interface inside the product
  - IP development/porting complexity - Effort to port into a new process node and potential schedule impact

Options

Option	Description
AIB	Intel's AIB using microbumps
Si Interposer	"Generic" silicon interposer using microbumps
BOW	Bunch of wires with standard bump pitch/packaging
BOW Turbo	Simultaneous bidirectional BOW
AQ Link	Aquantia proposal – Simultaneous bidirectional SERDES
Kandou	Kandou proposal – five signals on six wires
XSR	Optical XSR based PHY

- **We need the participation of experts in parallel interfaces, PCIe**

Relative Ratings for Parameters

	AIB	BOW Interposer	BOW	BOW Turbo	AQ Link	Kandou	XSR
Substrate	Interposer		Organic				
Type	Parallel				Serial		
Licensing	Open				RAND		Open
Piece cost	3	3	1	1	1	1	1
Operating cost*	(1)	(1)	(3)	(2)	(2)	(2)	(1)
Chiplet/product design NRE/schedule rick	2	2	1	2	2	2	2
Interface technology NRE/schedule risk	2	2	1	2	2	2	3

* Based on qualitative voting, reconciliation with FOM scores is still needed

Relative Importance of Parameters vs. Addressed Markets

	Sells value in the package	Sells value at the package		Sells value above the package		
Manufacturer	Chiplet vendor	Total BW < 5 Tbps	Total BW > 5 Tbps	Mobile	Edge	Data Center
Piece cost		4				
Operating cost		3				
Chiplet/product design NRE/schedule rick		1				
Interface technology NRE/schedule risk		2				

- **We need your feedback on your priorities for the market segment(s) in which you deliver products**

Thank You

Consume. Collaborate. Contribute.

Second Order Constraints – Likely Irrelevant

- Implementation complexity
- Number of unique power supplies
  - Limit interfaces to one high supply (1.5V?), one analog supply (0.9V?), and one digital supply (0.7V?)
- Ability to share power supplies
- Ability to negotiate to lower frequencies