



BoW: A Novel Interface for Inter-Chiplet Communication

ODSA Project Workshop
Sept 12, 2019

Ramin Farjadrad
Bapi Vinkakota
Mark Kuemerle

Bunch of Wires Update

- Who needs a new interface?
- What is the BoW and what's new with it?
- Where can it be applied?
- When can we get one?
- Why are we doing this?
- How can you help?

Who needs a BoW?

- Lots of people
- Yesterday (1990s): Everything was parallel on a board
- Today: Everything is serial on a board
- Tomorrow: It's a not so clear
 - Module to module
 - Serial still wins
 - Die to die on module
 - Serial interconnect -> maximize bandwidth per wire at the cost of latency and power
 - Die to die applications also need simple, low latency, low power interconnect
 - Package technologies have more wires than you think




What is the BoW?

- Source synch DDR interface - 16 single ended data and differential clock
- Performance Targets:
 - Throughput Efficiency → 100Gbps/mm - 1Tbps/mm (die edge)
 - Energy Efficiency → 0.5 pJ/bit - 1pJ/bit
 - Trace length → 1mm – 50mm
 - Low Latency < 5ns
- Small silicon area per port for dense integration to minimize area overhead
- Minimal analog/complex circuitry to offer easy/fast process porting
 - Limit the maximum baud rate of the interface → ~10G-16Gbaud
 - Common clock to replace CDR with DLL, Digital PLL to generate clock, etc.
- Single supply solution supporting wide V_{dd} range: 0.70V – 0.9V
 - Compatible with most existing SoC/ASIC in popular/available process nodes
- Optional ECC for ultra-low BER applications (BER<1E-30)

What's new with the Bunch of Wires Definition

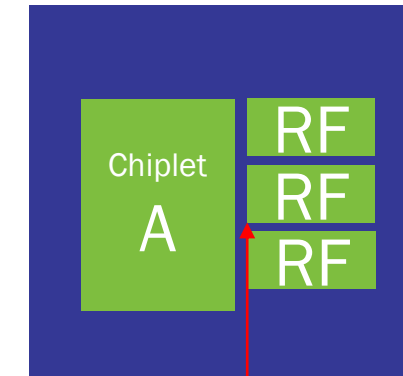
- Working toward specification release 0.7 for Amsterdam Regional Summit



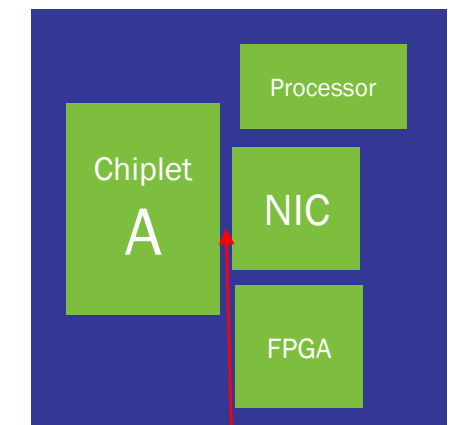
- Simplifying definition
 - Turbo mode moved to future expansion
 - Clarifying approach to bump definition
 - Pitch is not enforced, Order is → enables advanced packaging where needed
 - Power definition not enforced
 - Evaluating whether Fast and Base really need to be different definitions
- Channel definition
- Begging for your input and requirements (OK, not really new)

Where can BoW be applied?

- There are multiple proprietary versions of a simple parallel interface already being built
 - Open standard makes sense to converge design points
 - Must meet needs of a broad spectrum of applications
- A simple parallel interface at relatively low datarate enables chips from older technologies to be integrated into the system
 - Critical for RF component integration
- Expandability of the same pins at higher datarates allows for chiplets to meet multiple application scenarios
 - Eg. Streaming data with several RF chips
 - Same die switching accelerator data with NIC, FPGA



Hundreds of Gbps



Thousands of Gbps

When can we get one?

- Targeting release of spec end of Sept 2019
- ODSA strongly supports
 - Open implementations of BoW to spread the joy
 - Help from fabricators to prototype the IP
 - Please contact us if you'd like to help

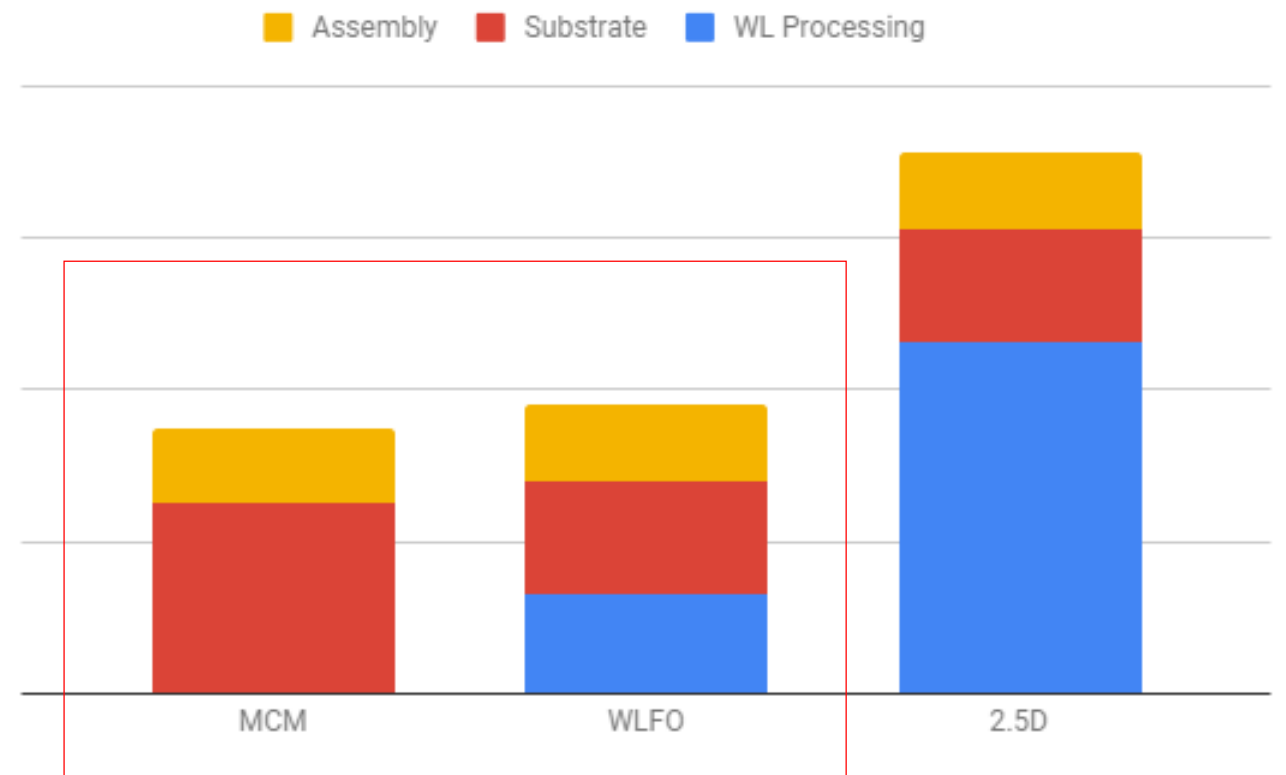
Why are we doing this?

- Because its simple and affordable to implement
 - Low bandwidth interconnect can use very simple IP and package
- Because anyone can use it
 - BoW definition is very simple with off the shelf components
 - More bandwidth can be enabled with more complexity
- Having an open, common interface enables interoperability between chiplets
 - Increases the types of solutions available

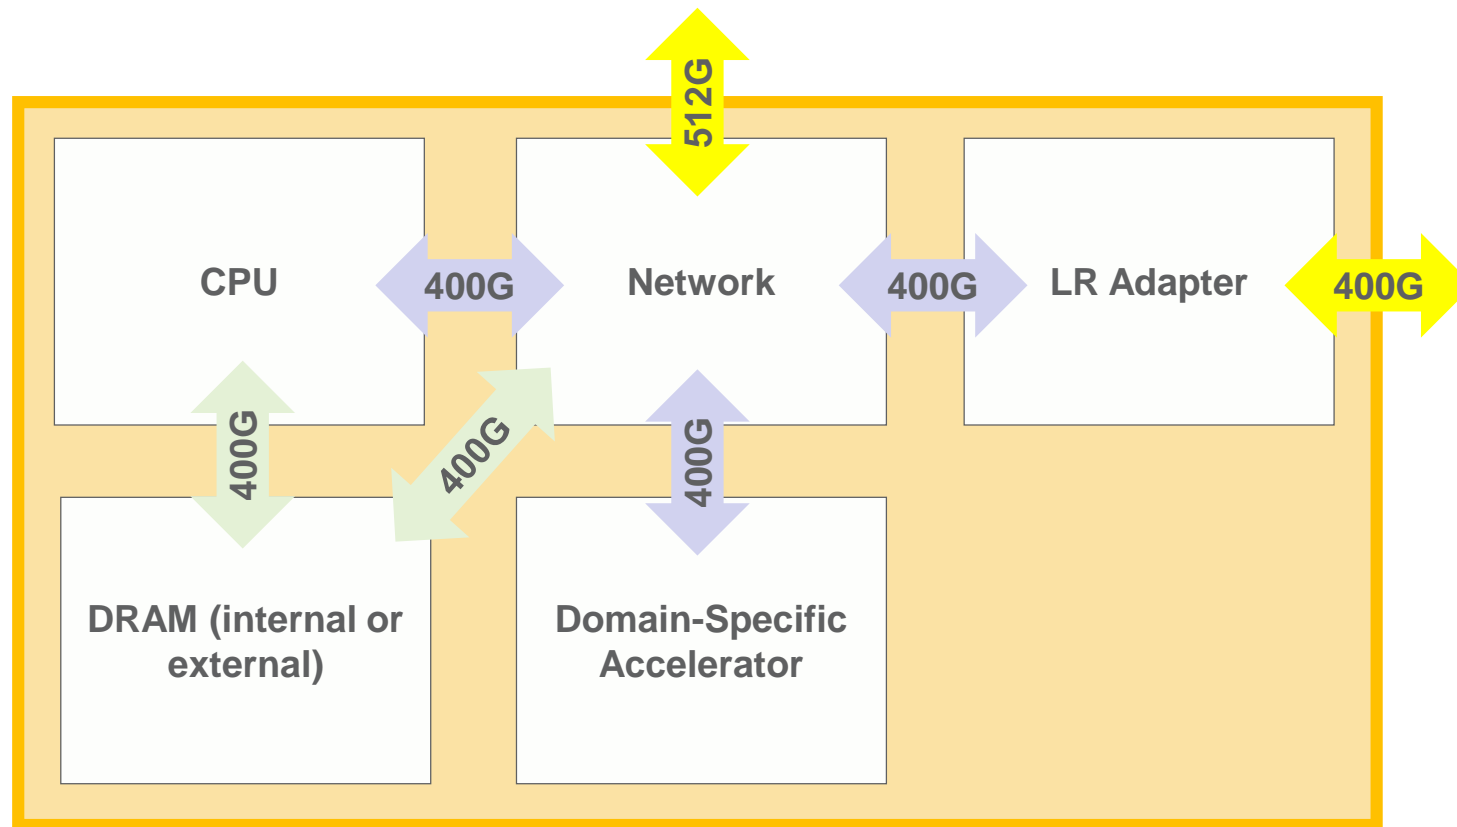
Why build an interface compatible with standard laminates?

- Advanced packaging technology can be expensive
- TSV based technologies nearly double package cost
- Fanout technologies might mitigate some of cost increase by reducing overall substrate cost

WL Processing, Substrate and Assembly



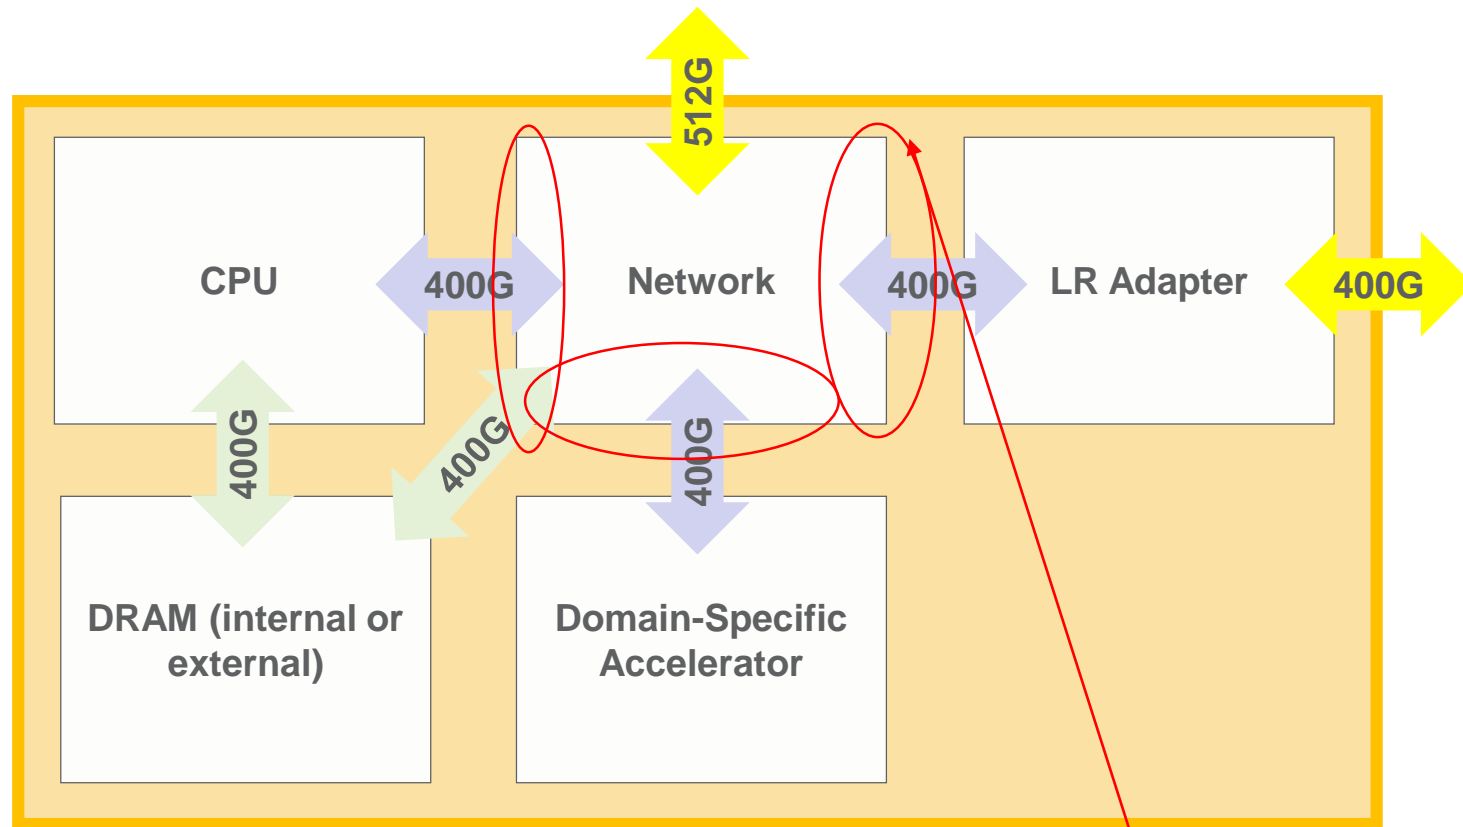
BoW for Board to Package Aggregation (reference)



- Open Domain-Specific Architecture – OCP project: <https://www.opencompute.org/wiki/Server/ODSA>
- Reference accelerator architecture
- Today implemented as a board → Will be implemented as a single package.
- PCIe on blue links, today on PCIe PHY → Ported to BoW in a single package.

Interface parameters			
PHY	PCIe	BoW Basic	BoW Fast
Trans. Protocol	PCIe4	PCIe4	PCIe4
pJ/bit	7.5	~0.6	~0.7
Cost of 512 Gbps interface (2 x 16 PCIe Gen4 lanes, Tx=Rx=512Gbps)			
Power	3.84W	0.3W	0.35W
Bump count		416	104
Area sq mm.	3.9	5.1	1.8
Impact on PoC (3 x 512 Gbps interfaces)			
Total power	11.5W	0.9W	1.05W

BoW for Board to Package Aggregation




~1200 Gbps bidi in >12 mm of chip edge, 200Gbps unidirectional

Why support a range of bump pitches?

BoW meets bandwidth needs of most accelerators with low package costs

BoW can be used with more advanced packaging to solve the most demanding applications



3 rows of BoW Base are enough for most accelerator bandwidth requirements

All options for BoW Fast exceed bandwidth requirements

BoW Fast with Fanout technology nearly as good as Base 2.5D at a fraction of the package cost

How Can You Help?

- Develop areas of the specification that need further definition
 - Initialization & Calibration
 - Test
- Develop & Share implementations
- Support implementation development

Thank You!