

Open-Domain Specific Architecture: An Overview

Bapi Vinnakota

Sub-project lead: ODSA, Open Compute Project
(representing an active community of contributors
from over 30 companies)
Broadcom Inc.

Presented at:
MEPTEC-IMAPS Semiconductor
Speaker Series October 14, 2020



OPEN
COMMUNITY®



Outline

- Motivation: Chiplets, ODSA
- Overview: Community, charter
- Review: Significant results, 2020 plans
- How to participate, call to action

What this is not - a technical deep dive.
Pointers to detail through the talk



OPEN DOMAIN
SPECIFIC
ARCHITECTURE



SERVER



Chipllets: More Attention

DESIGNLINES | MEMORY DESIGNLINE

Chiplet Uptake Creates Demand for Best Practices

By Gary Hilson 07.08.2020 0

Share Post [Share on Facebook](#)

EE Times
07/08/20

TORONTO — Chiplets are a great example of a solution that's been around for a while but is quickly finding more problems to solve.

With Moore's Law now 55 years old and pace of semiconductor manufacturing advancement decelerating, chiplets offer an approach to semiconductor design and integration that hold the promise of speeding up things up again. Recent research released by Omdia forecasts the global market for processor microchips that use chiplets in their manufacturing process to hit \$5.8 billion in 2024, a significant jump from \$645 million in 2018.


AMD says chiplet design can cut costs by more than half

AMD has consistently beaten Intel in cost-per-core, but did AMD really need to pursue a chiplet design to make Zen 2 so affordable? New slides from a recent talk at ISSCC show exactly how much the company saved with this approach, and the results are very impressive.

Jonathan Hayhurst, 02/28/2020

Notebookcheck.net
02/28/20

Since the launch of its 2019 Zen 2 processor, AMD's chiplet design is at its most consistently undercut Intel from a dollars-per-core perspective. AMD's value oriented approach can be attributed to many factors, but none of them are more relevant than the novel "chiplet" design AMD uses in manufacturing their latest Zen 2 chips. The idea is to take several smaller dies manufactured on different processes, and put them together on one package to improve yields and thereby reduce costs. But how much can they reduce them by how much? Well by more than half in some cases.



Chiplet design is at its most extreme in AMD's 64-core EPYC CPUs (Image source: Wired)

The Next Advanced Packages

752 Shares

f 122

22

in 467

<

New approaches aim for better performance, more flexibility — and for some, lower cost.

JUNE 18TH, 2020 - BY: MARK LAPINSKY

Semi Engineering
06/18/20



Packaging houses are redefining chiplet packages, paving the way toward new and innovative system-level chip designs.

These packages include new versions of [2.5D/3D](#) technologies, [chiplets](#), [fan-out](#) and even wafer-scale packaging. A given package type may include several variations. For example, vendors are developing fan-out packages and panels. One is combining fan-out

3 Ways Chiplets Are Remaking Processors

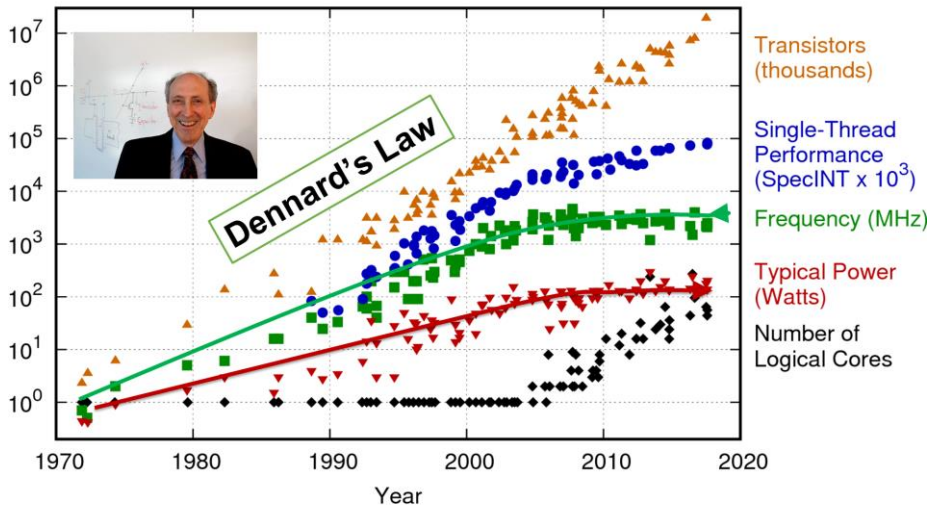
AMD and Intel are leaning on chiplets to boost performance; CEA-Leti shows just how far the approach can go

IEEE Spectrum
04/29/20



Slower Scaling, Higher Costs


42 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
 New plot and data collected for 2010-2017 by K. Rupp data source: <https://goo.gl/bb6wZW>

Using 'more than Moore' technologies can significantly cut the cost of an ASIC

TSMC's available processes: May 2019



Using mature technologies cuts the mask set costs significantly: 2018 vs cost when each technology was introduced



The economics of ASICs: At What Point Does a Custom SoC Become Viable?: Electronic Design, 7/15/19

Domain-Specific Accelerators

Augment general-purpose CPUs

Programmable silicon with a data path optimized for a compute-intensive application

Neural-net training, inferencing, video, encryption, crypto....

<https://engineering.fb.com/data-center-engineering/accelerating-infrastructure/>

<https://cloud.google.com/blog/products/gcp/google-supercharges-machine-learning-tasks-with-custom-chip>

<https://www.computeexpresslink.org/>

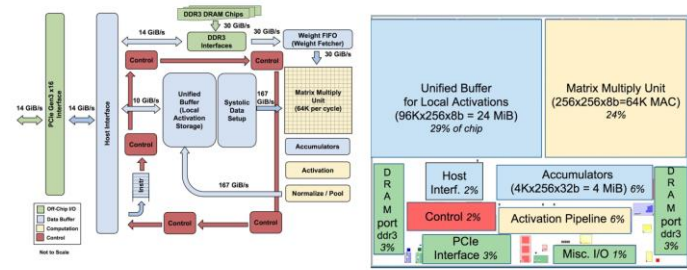



Figure 1. TPU Block Diagram. The main computation part is the yellow Matrix Multiply unit in the upper right hand corner. Its inputs are the blue Weight FIFO and the blue Unified Buffer (UB) and its output is the blue Accumulators (Acc). The yellow Activation Unit performs the nonlinear functions on the Acc, which go to the UB.

Figure 2. Floor Plan of TPU die. The shading follows Figure 1. The light (blue) data buffers are 37% of the die, the light (yellow) compute is 30%, the medium (green) I/O is 10%, and the dark (red) control is just 2%. Control is much larger (and much more difficult to design) in a CPU or GPU.




Compute Express Link™: The Breakthrough CPU-to-Device Interconnect

About CXL™

Gordon Moore's "Other" Observation

"It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected."

Electronics, volume 38, number 8, April 19, 1965



Chiplet-Based Products

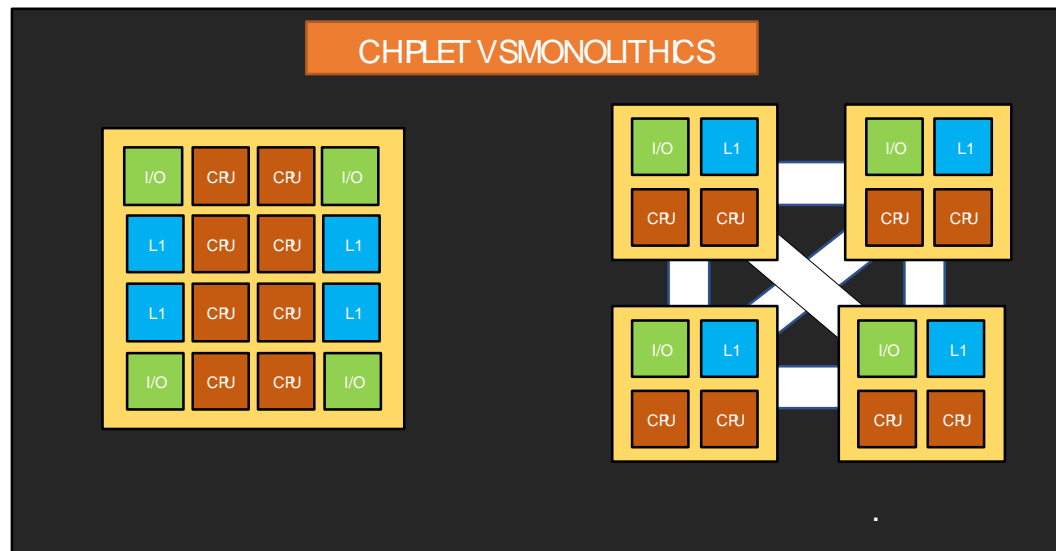
Heterogeneous integration: modular design integrating die from multiple process nodes.

Reduce design, manufacturing cost, preserve near-monolithic performance (ODSA white paper)

Need an energy-efficient PHY and logical die-to-die interface

Today: Proprietary D2D interfaces, single-vendor multi-chiplet products

PROVEN IN EXISTING BUSINESS MODELS



From Kevin Drucker, Facebook
Talk at Broadcom

[L. Su IEDM' 17]

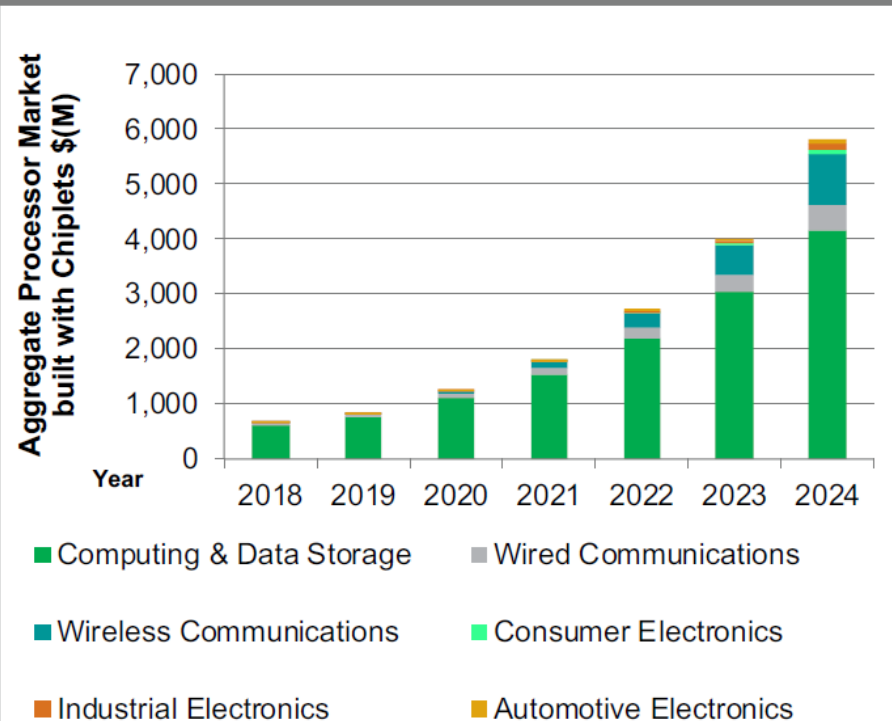
Market Forecast by Omdia

Chiplet market expected to grow to ~\$6B by 2024 and to double by 2030

“Hyperscalers, cloud and communications service providers all have incentives to see the computing and data storage market become more resistant to the eminently increasing expense associated with the end of Moore’s Law.”

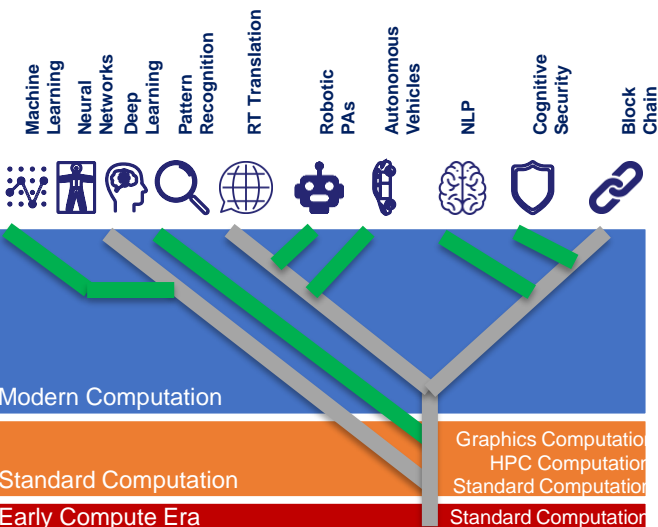
“Compute & Storage Space dominates early growth. There is already significant chiplet technologies developed as proprietary solutions.”

Overall Chiplet SAM by Market Segment




ODSA: Accelerators and Chiplets

Domain-specific architectures (DSAs) to accelerate targeted compute-intensive workloads.



AI/ML/data workload explosion needs DSAs


Dharmesh Jani, Facebook – ODSA Workshop, Regional Summit, Amsterdam, Sep. 2019



OPEN DOMAIN
SPECIFIC
ARCHITECTURE

DSAs built using chiplets with open standard D2D interfaces

Chiplet: Die designed to be used with other die in a package, usually with proprietary interfaces.




IBM Power 9: potential modularity

Jeff Stuechli, Josh Friedrich, IBM – ODSA Workshop, IBM, San Jose, Sep. 2019

Attendees and Participants:

Attendance and/or participation do not imply corporate endorsement



Semi Vendors
IP providers, EDA
Service providers

Tools, Manufacture,
Design, Test,
Integration

Systems vendors,
End users, ISVs, Service
Providers

A growing community

ODSA Workstreams

Attendance and/or participation do not imply corporate endorsement








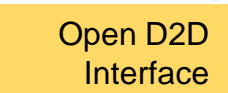






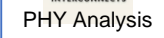
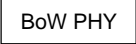










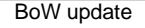
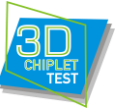






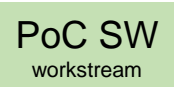


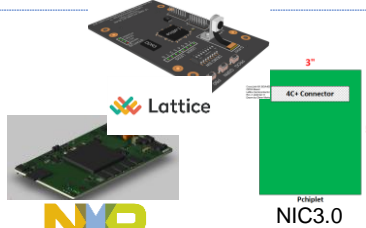
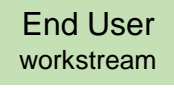

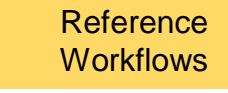

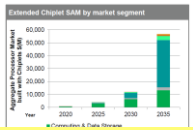
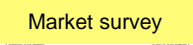
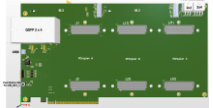

Each group meets weekly, details at <https://www.opencompute.org/wiki/Server/ODSA>

Workstream	Leader	Participants	Objective
PHY Layer	Robert Wang 	    	PCIe PIPE adapter
Bunch of Wires	Mark Kuemerle 	     	Low cost D2D PHY
CDX	Jawad Nasrullah 	    	Chiplet design exchange
Business	Sam Fuller 	  	Chiplet workflowP
PoC hardware	JP Balachandran 	   	PoC board design
PoC software	Kevin Drucker 	   	Application/Infra software
Link layer	Open	    	ODSA Stack
OpenHBI	Kenneth Ma 	    	High perf D2D PHY
End user	Dharmesh Jani 	  	End user input to the ODSA

Sub-project lead: Bapi Vinnakota 

ODSA Progress

details at <https://www.opencompute.org/wiki/Server/ODSA>

<p>Workshop   </p> <p>Announcement </p> <p> CDX workstream</p> <p> PHY workstream</p> <p> LINK workstream</p> <p> Open D2D Interface</p> <p> White paper</p>	<p> </p> <p> </p> <p> BoW 0.7</p> <p> PHY Analysis</p> <p> BoW PHY</p> <p> BoW workstream</p>	<p>  </p> <p> OHBI workstream</p> <p> </p> <p> PIPE 0.5</p>	<p></p> <p> ODSA Intro</p> <p> BoW update</p> <p> </p> <p> </p>
<p> Reference Designs</p> <p> PoC workstream</p>	<p> PoC design</p> <p> PoC SW workstream</p>	<p> </p> <p></p>	<p> End User workstream</p> <p></p>
<p> Reference Workflows</p> <p> BIZ workstream</p>	<p></p> <p> Market survey</p>	<p></p> <p> PoC pchilets, mara</p>	


1H, 2019

2H, 2019

1H, 2020


2H, 2020

D2D Interface: Use Cases




Standardized motherboard interfaces enable the PC ecosystem

Package Level Integration



Standardized chiplet interfaces enable a package-level integration ecosystem

SOC Disaggregation



Standardized SOC interfaces (AMBA/AXI) enable foundry ecosystem

From Board to Package

Benefits

Smaller form factor
Higher bandwidth, lower latency
Power efficiency

Challenges

Business models
Known good die,
Thermal limits

To Package from SOC

Benefits


IP portability/suitability
Potential lower NRE cost & TTM
Address reticle size limits, yield

Challenges


Form factor, bandwidth, latency,
Silicon area/power overhead
Manufacturing cost

R. Nagisetty, Intel – ODSA Workshop, Intel, June 2019
J. Friedrich, IBM – OSA Workshop, Sep 2019
R. Cheema, Socionext – ODSA weekly meeting, July 2020


D2D interface: Packaging Options




- Substrate interconnect
- Low/med density D2D
- > 15um line/space
- > 125um bump pitch
- > 50um comp spacing
- Bare/packaged parts
- Chip last process




- Substrate interconnect
- Low/med density D2D
- > 10um line/space
- > 125um bump pitch
- > 50um D2D spacing
- Chip last process



- RDL interconnect
- Med/high density D2D
- > 1um line/space
- > 40um bump pitch
- 100um D2D spacing
- Chip first/last
- FC to organic subs.
- No TSV – lower loss



- Si interconnect
- High density D2D
- < 0.5um line/space
- > 40um bump pitch
- High bandwidth
- > 50um D2D spacing
- Chip first/last
- Power limitation

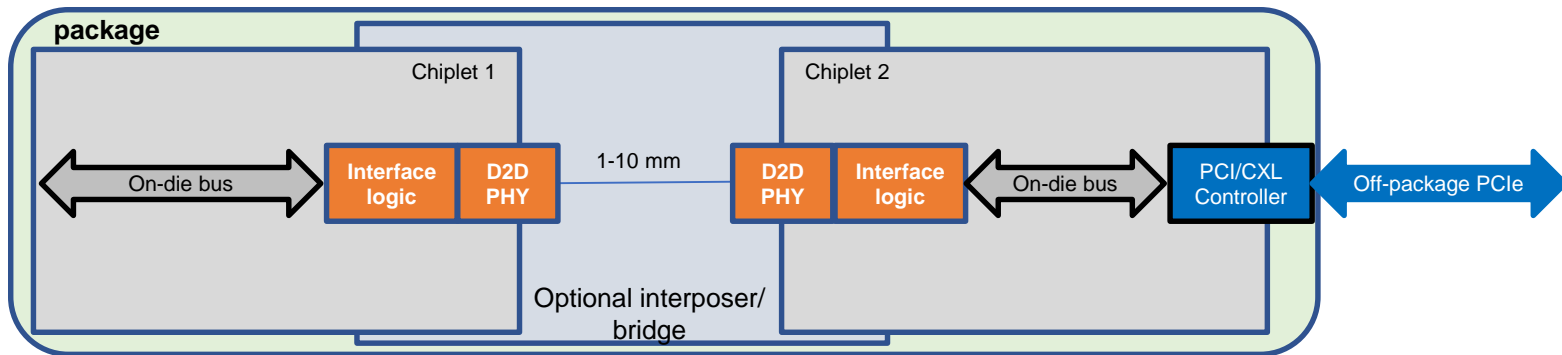


Development

- Si interconnect
- High density D2D
- Foundry line/space
- < 40um bump pitch
- High bandwidth
- Shortest interconnect
- FC or hybrid bond
- Custom die designs

Increasing Bandwidth & Cost

Open D2D Interface Requirements



PHY: Enable heterogeneous integration

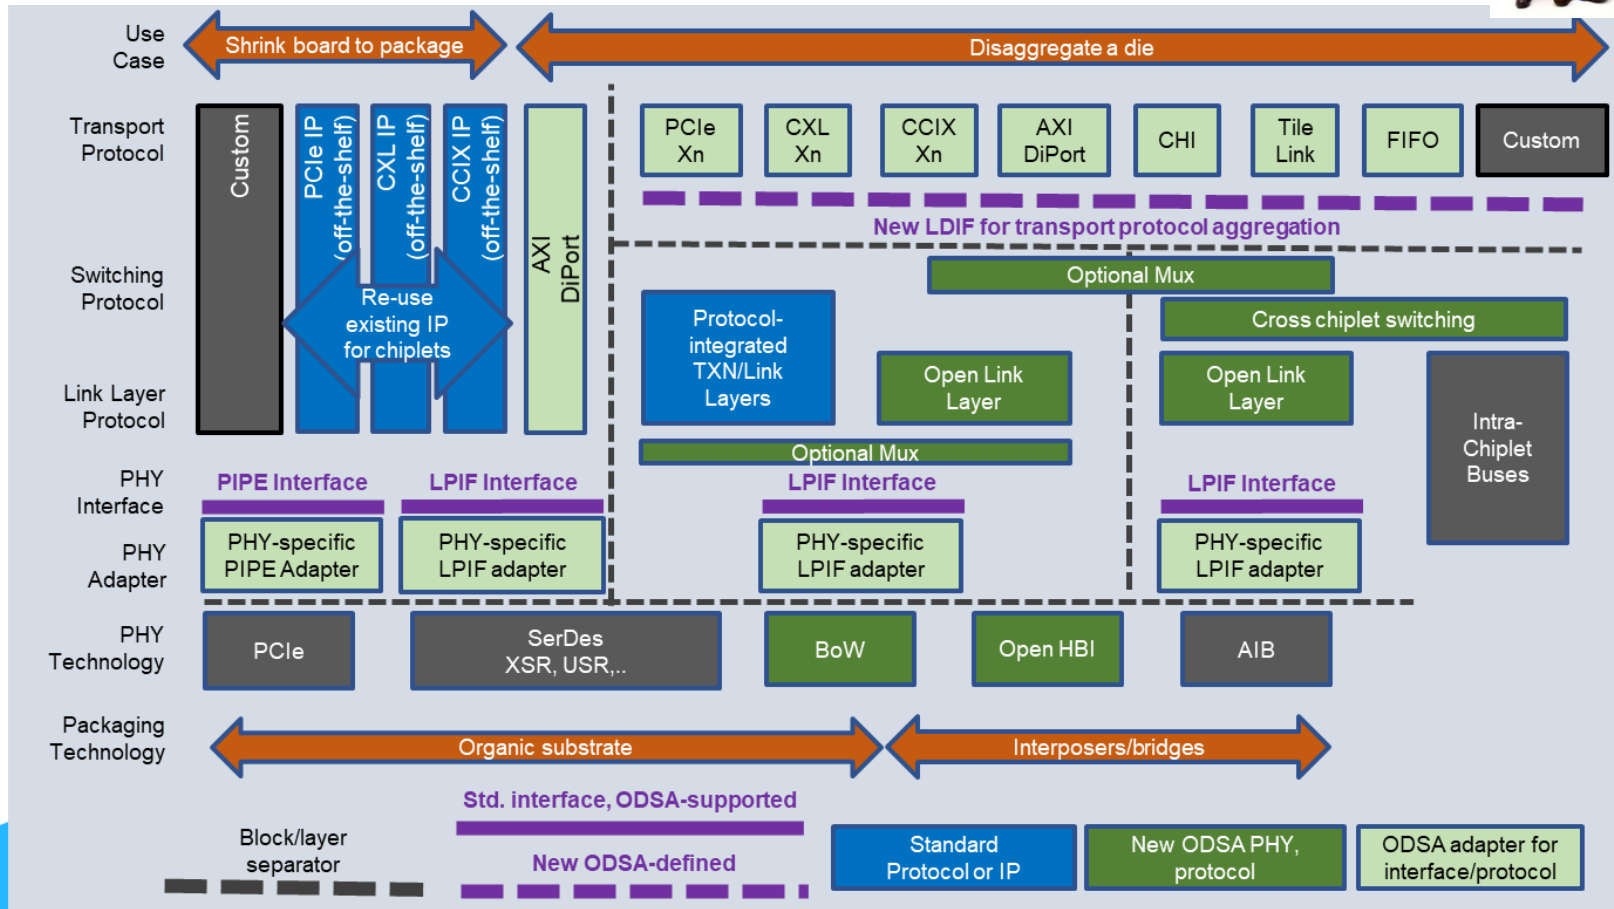
- Low power PHY – 0.5 pJ/bit
- Usable across packaging technologies, process nodes

Logic: Underlay for transaction protocols for two use cases

- PCIe, CXL, CCIX to shrink a board to a package.
- AXI, CHI, Proprietary, TileLink buses to disaggregate a design

ODSA Stack:

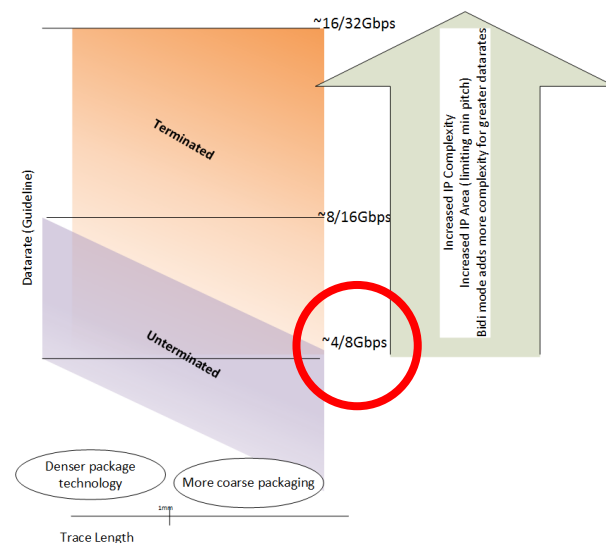
K. Drucker et al – ODSA, Hot Interconnect, 2020



Bunch of Wires PHY

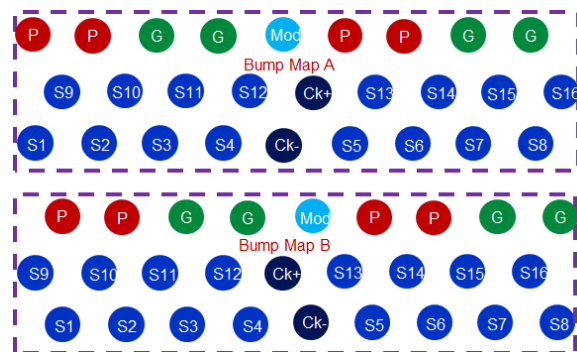
Open D2D PHY:

- Simple clock-forward base parallel PHY – 4-8 Gbps/wire, < 5 ns latency, 0.75V
- Supports process nodes from 3nm to 65nm to enable heterogeneous designs.



Only PHY to offer a graceful cost-performance trade-off


Type	Design	Packaging	Power	Performance
Simple design Simple package	Base clock forward parallel	Simple organic C-4	~0.5 pJ/bit	0.28 Tbps/mm
More design Simple package	Destination terminated	Simple organic C-4	~0.6 pJ/bit	1.3 Tbps/mm
Simple design More package	Base clock forward parallel	WLFO/ interposer	< 0.5 pJ/bit	1.8 Tbps/mm



Base unit: 16-bit slice
~4xPCIe Gen 4 lanes

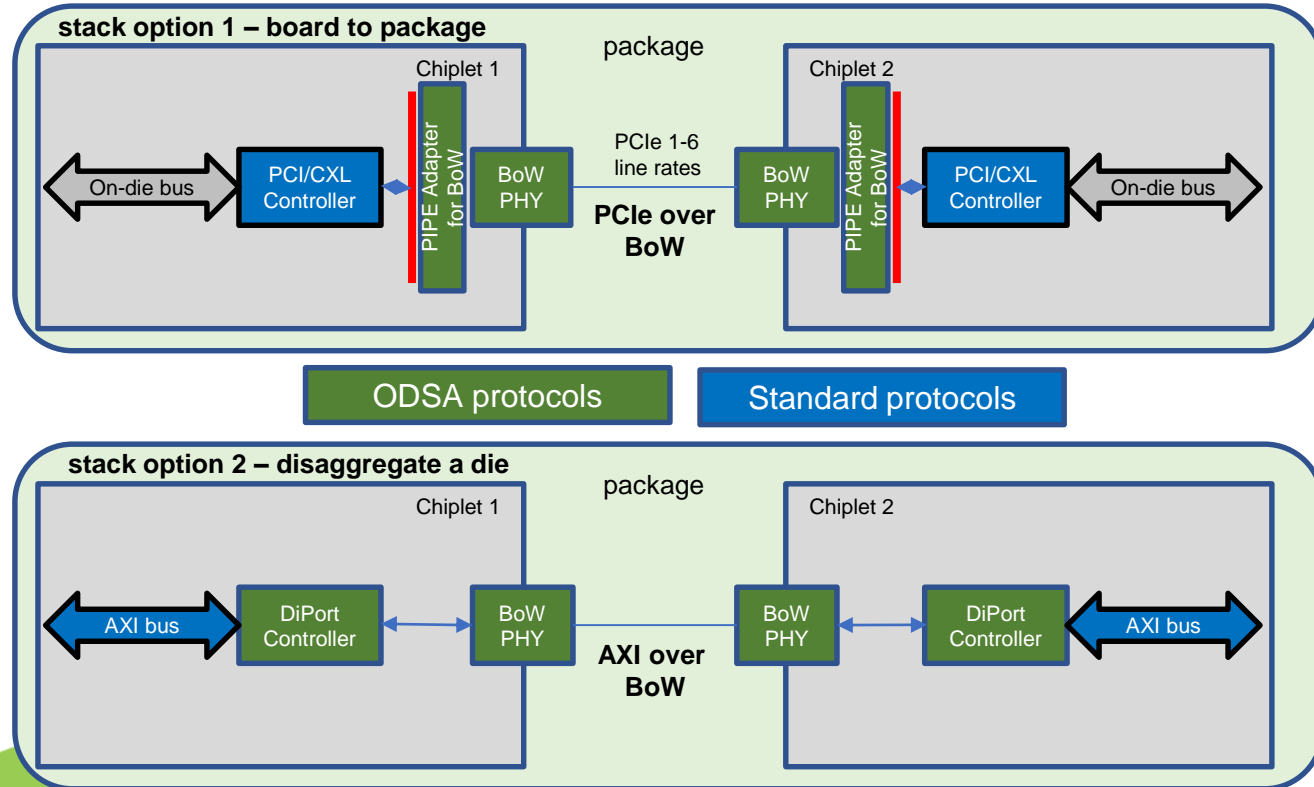
2020: ODSA PHY/Logic D2D Interface

K. Drucker et al – ODSA, Hot Interconnect, 2020

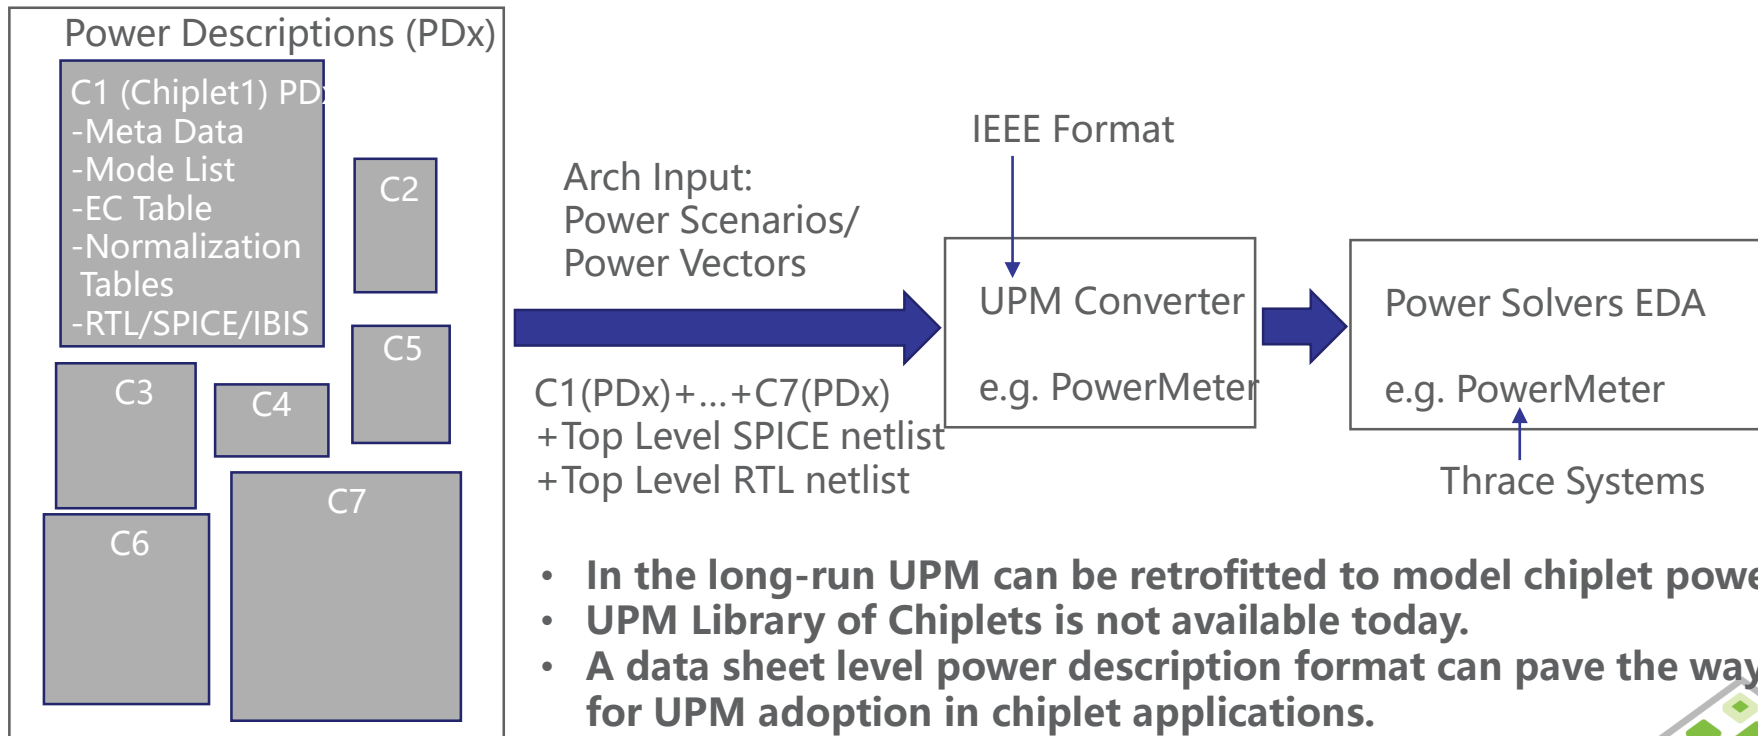


Port the most common system (PCIe/CXL) and SoC (AXI) transaction to chiplets.

- PCIe/CXL over BoW through standard PIPE interface
- AXI over BoW with DiPort contributed by NXP



Multichiplet Design Workflow: Proposed Chiplet Power Description

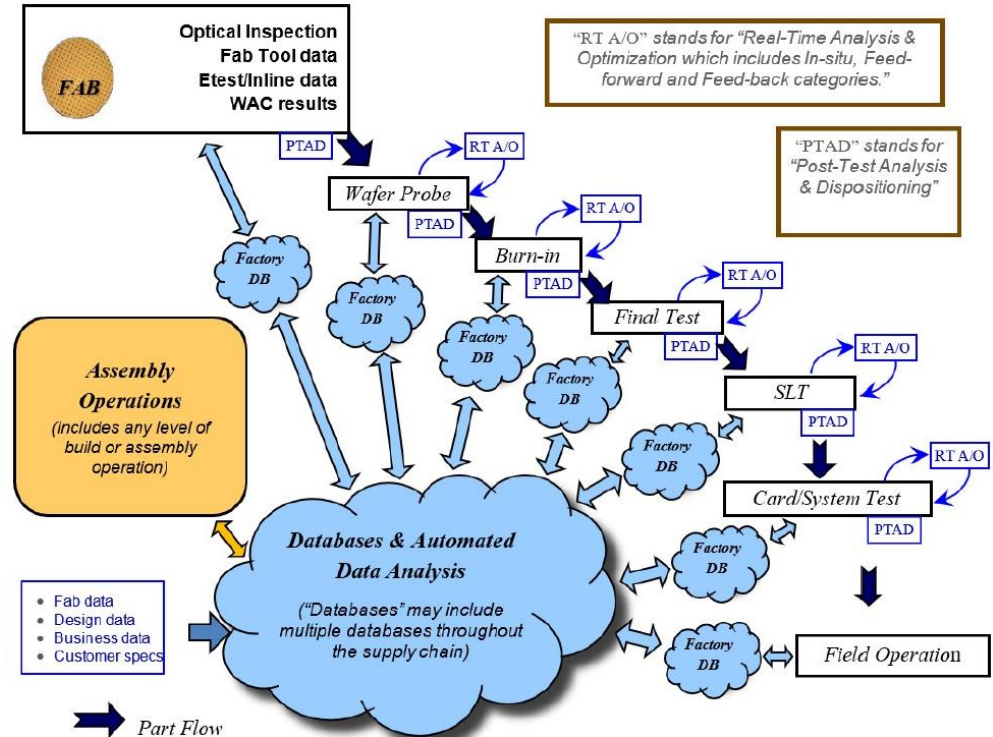


- In the long-run UPM can be retrofitted to model chiplet power.
- UPM Library of Chiplets is not available today.
- A data sheet level power description format can pave the way for UPM adoption in chiplet applications.

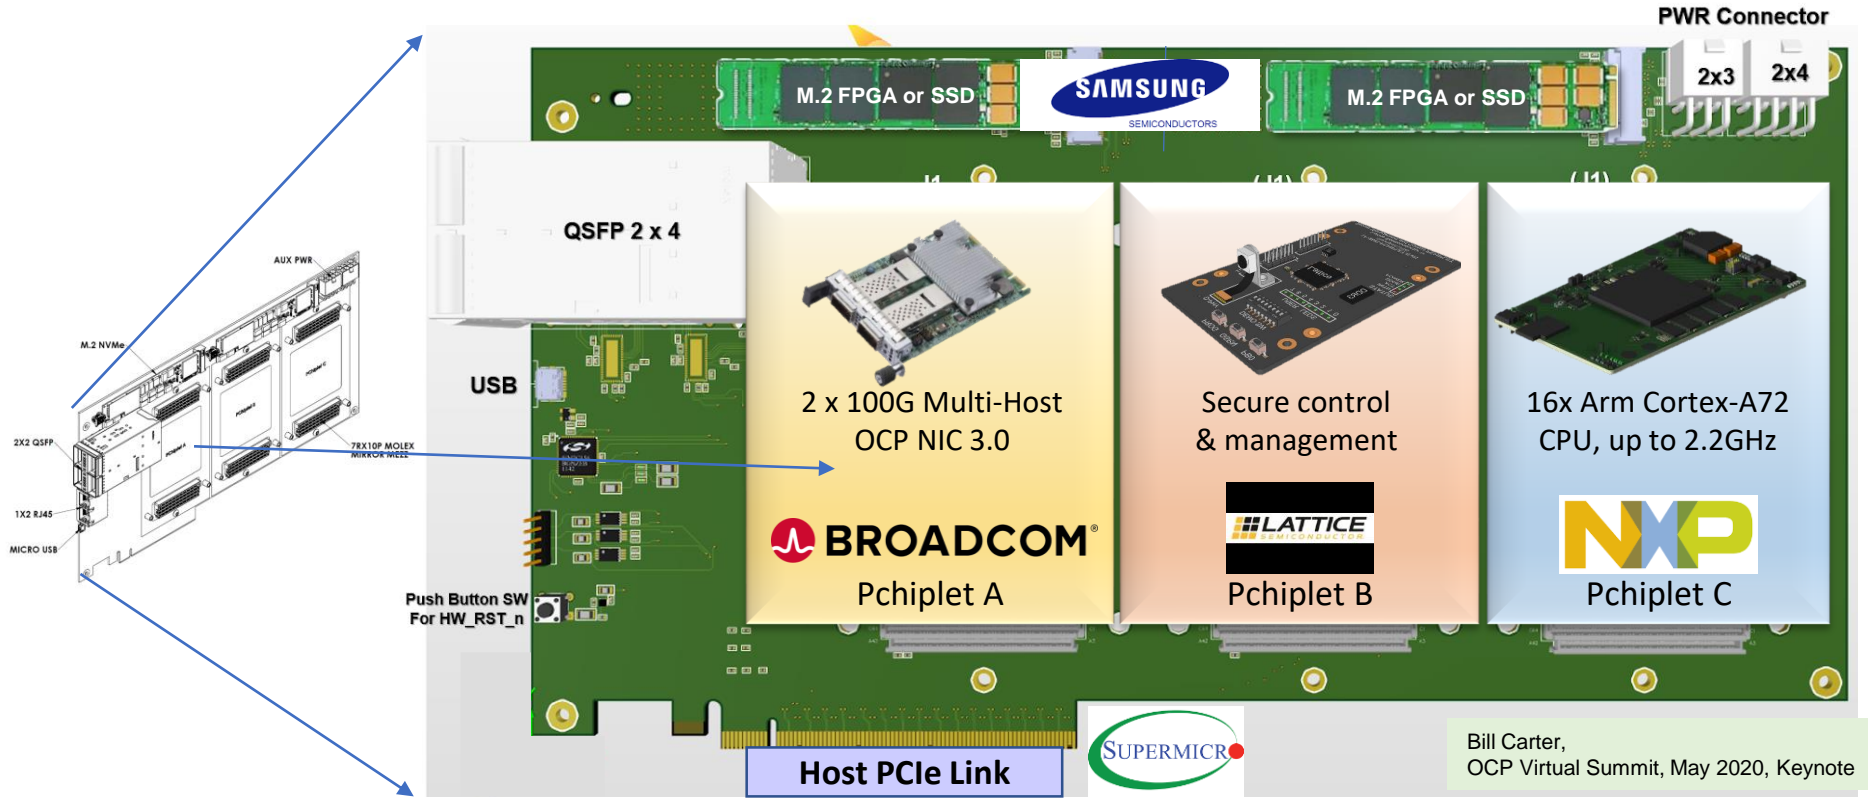


Multichiplet Test Workflow: Data Exchange

- Improve feedback loops for diagnosis and low Cost of Test
- Tests could be moved in the flow.
- Enable database across chiplet vendors and product developers.
- Enable vendors to respond to specific yield requirements
- preserving business-confidential information



ODSA Workflow PoC Kit

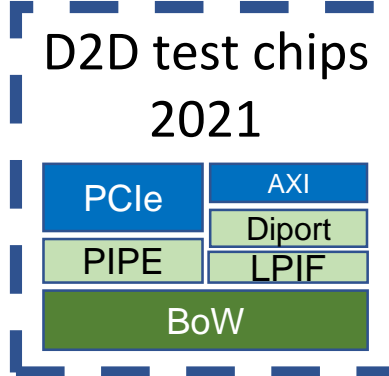

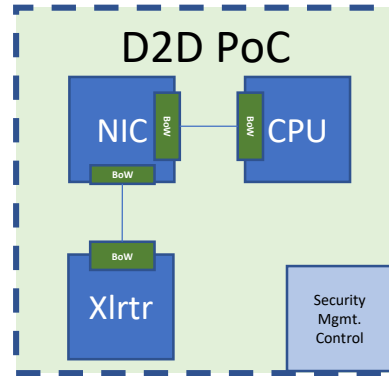
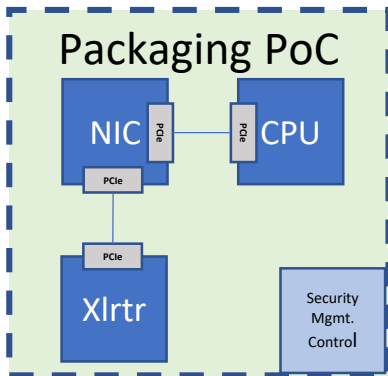



Design your own Pchiplet, develop an application

PoC Roadmap

Chips to Chiplets
Open Workflow

D2D optimized
Chiplets Workflow



— Current ODSA activities

- - - Roadmap



For More Information

[ODSA Wiki](#): All workshops, minutes of weekly calls, workstreams (open access) – talks on business model, co-packaged optics, design, packaging, memory chiplets, power modeling, XSR...

Specification proposals:

Bunch of Wires [GitHub repo](#) (open access)

[PIPE adapter](#) (PCIe over D2D), [DiPort](#) (AXI over D2D) (open, but need to request access)

[ODSA PoC Demo](#) (open access), [ODSA PoC Implementation Specification](#) (open, but need to request access)

In-flight

LPIF, LPIF' proposal

ODSA PoC SW

Open HBI specification (needs CLA)

Technical Papers/Talks

[ODSA white paper](#), ODSA Wiki

R. Farjadrad, M. Kuemerle, B. Vinnakota, “A Bunch-of-Wires (BoW) Interface for Interchiplet Communication”, IEEE Micro, 2020

G. Taylor, R. Farjadrad, B. Vinnakota, “High Capacity On-Package Physical Link Considerations”, Hot Interconnects, Aug. 2019

D. Jani, “Musings on Domain Specific Accelerators, Open Compute Project and Cambrian Explosion”, LinkedIn

M. Hutner, et al “Test Challenges in a Chiplet Marketplace”, VLSI Test Symposium, Apr. 2020

B. Vinnakota, “The Open Domain-Specific Architecture: An Introduction”, Design Automation Conference, July. 2020

S. Ardalan et al, “Bunch of Wires PHY: An Open D2D Interface”, Hot Interconnect 2020

K. Drucker et al, “The Open Domain-Specific Architecture”, Hot Interconnect 2020

S. Ardalan et al, “BoW Interface: Interchiplet Link Testing and Loopback”, 7th Int. Work. on 3D and Chiplet Test, Nov 2020




ODSA at the OCP Tech Summit

Area	Title	Speakers	Company
General	ODSA Status Update	Bapi Vinnakota	Broadcom
	End user panel	D. Jani, D. Xu, R. Mittal, M. Chowdhry	Facebook, Alibaba, Google, Microsoft
Open D2D interface	The Bunch of Wires 1.0 release	Ken Poulton	Keysight
	The ODSA Pipe Adapter	Michael Spear	IBM
	An Open Chiplet Link Layer	Arthur Marris	Cadence
	Open HBI introduction	Kenneth Ma	Xilinx
	D2D PHY Panel	M. Kuemerle, U. Sjöström, A. He, V. Kugel	Marvell, Ericsson, Google, Juniper
	D2D PHY Comparison	Bapi Vinnakota, Shahab Ardalan	Broadcom, Ayar
Reference designs	Bunch of Wires Test Chip	Suresh Subramanian	Apex
	NXP PoC Pchiplet	Sam Fuller	NXP
	Lattice PoC Pchiplet	Marshall Goldberg	Lattice
	ODSA PoC update	Jayaprakash Balachandran	Cisco
Reference workflows	Chiplet Design Exchange update	Jawad Nasrullah	zGlue
	Packaging for chiplets	Agarwal, Heung, Kelly, Chen, Tzou	Facebook, JCET, Amkor, ASE, TSMC
Company talks	Chiplet optics	Shahab Ardalan	Ayar Labs
	Chiplet packaging	Eelco Bergman	ASE

Please Help, Join Us!

- Join a work stream, each meets weekly
- Help with the PoC, software, use case dev, Q4 demo
- Review, help complete documents in flight
- Need packaging and test definition and work streams
- Make chiplets with, IP for, the open ODSA stack
- <https://www.opencompute.org/wiki/Server/ODSA>



OPEN DOMAIN
SPECIFIC
ARCHITECTURE



SERVER

