The Open Domain-Specific Architecture (ODSA)

Bapi Vinnakota

Representing an open community in the OCP

Broadcom Inc

Bio

Bapi Vinnakota is a System Architect with Broadcom. After a Ph.D. at Princeton, he taught at the University of Minnesota, where he received an NSF CAREER and IBM Faculty Development Awards.

He joined Intel through an acquisition and was an architect of a VoIP flow processor, worked in networking technology and incubated a networking SaaS product. At Netronome, he created and ran open-nfp.org, a service for research in networking.

He leads the Open Domain-Specific Architecture sub-project, in the Open Compute Project. The ODSA has active volunteers from over 30 companies and aims to define an open chiplet marketplace.

Overview

- ODSA focus: accelerators and chiplets
- ODSA motivation, charter, community
- Deep dive into the ODSA interface
- How to participate in the ODSA

ODSA: Accelerators and Chiplets

Domain-specific architectures (DSAs) to accelerate targeted compute-intensive workloads[1].

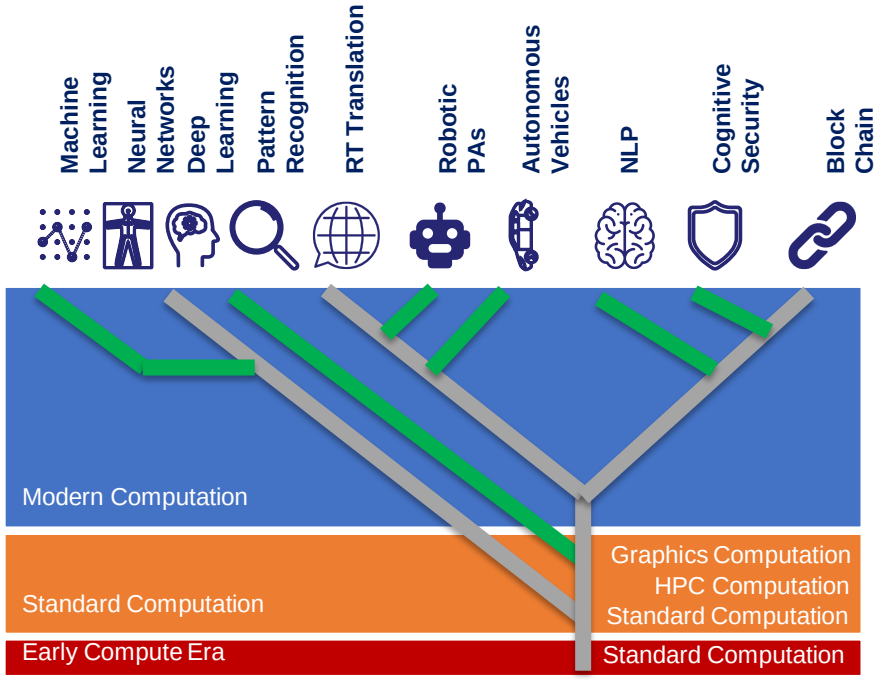

AI/ML/data workload explosion needs DSAs

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[2].

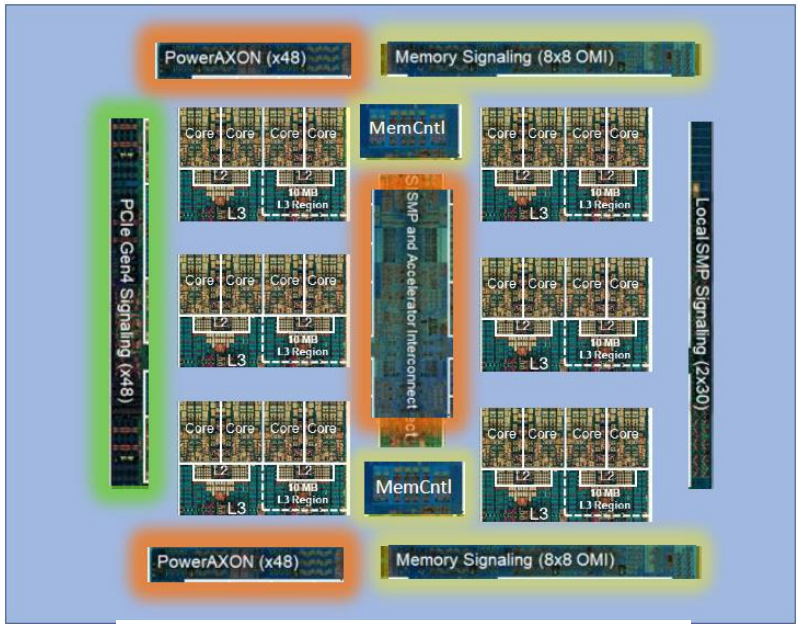

IBM Power 9: potential modularity

1. Dharmesh Jani, Facebook – ODSA Workshop Regional Summit, Amsterdam, Sep. 2019
2. Jeff Stuechli, Josh Friedrich, IBM – ODSA Workshop, IBM, San Jose, Sep. 2019

Chiplet-Based Products

- Heterogeneous integration: modular design across multiple die from multiple process nodes.
- Reduce design, manufacturing cost, choose optimal node for function, 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
- Tomorrow: An open D2D interface, multi-vendor multi-chiplet products

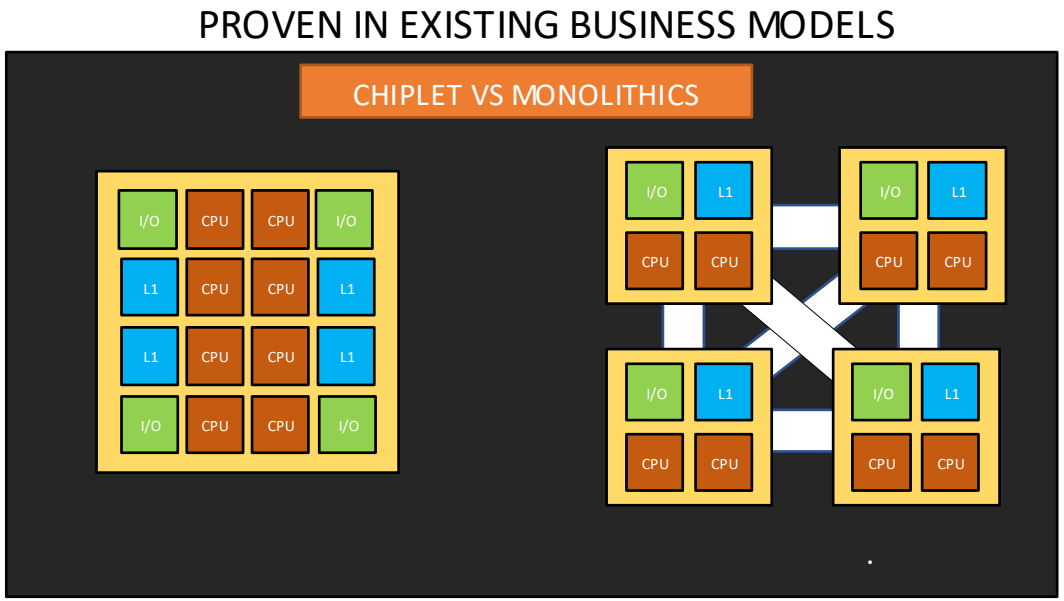

From Kevin Drucker, Facebook
Talk at Broadcom

[L. Su, IEDM'17]

D2D Interface and Packaging for Chiplets

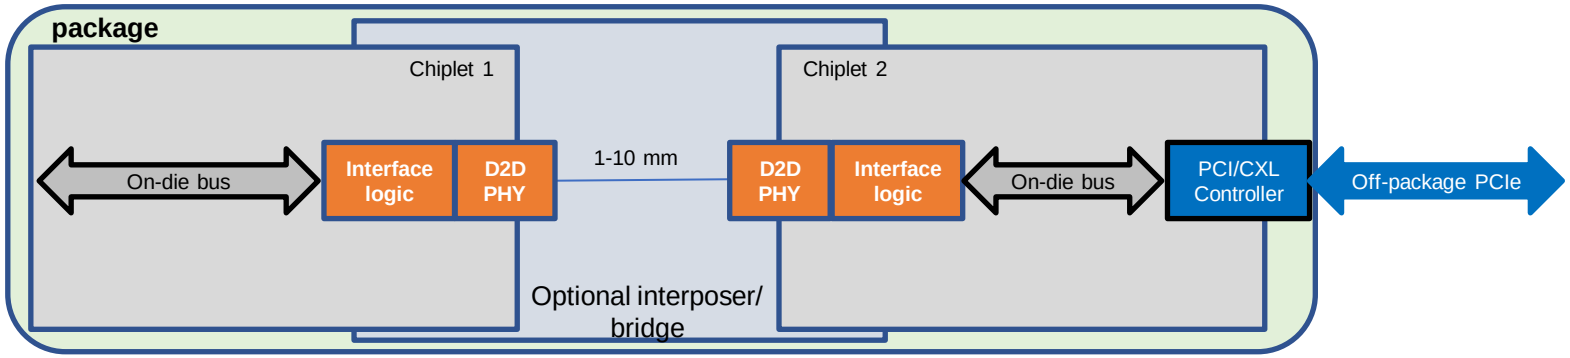

- Die need **PHYsical and logical connectivity** for modular design (ODSA white paper)
  - D2D PHYs 10x more energy efficient than off-package protocols, less than on-die buses
  - DARPA, Taylor et al (Hot Interconnect, 2019)
- Workflow for multiple chiplets in a package
  - Many packaging options[1]: organic laminate, wafer-level fanout, silicon interposer, silicon bridge
  - Need to use known good die in a package for acceptable product yield
- Use cases drive connectivity and packaging.
  - Two use cases (Intel/IBM talks June/Dec ODSA workshops):
  - Shrink a board to a package, Disaggregate a die

1. ASE, March 27th talk at the ODSA wiki

ODSA Charter

Open D2D Interface
Reduce barrier to interoperation

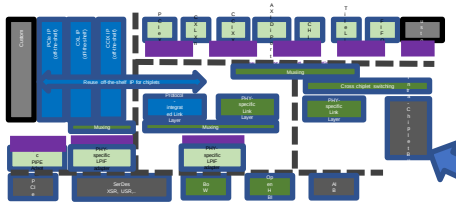

Reference Designs
Starting point for new designs

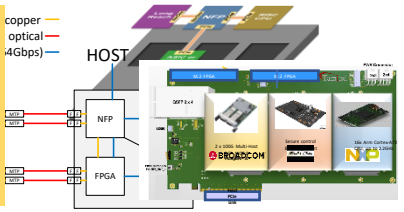

Reference Workflows
Reusable, open practices

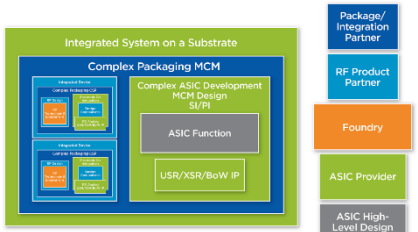

ODSA Activities

Chiplet Marketplace

Integrate best-in-class chiplets from multiple vendors through open interfaces

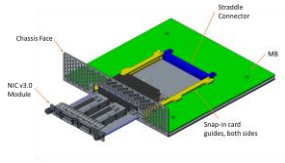

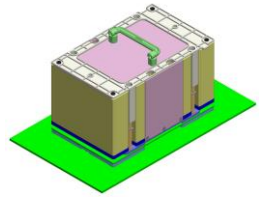

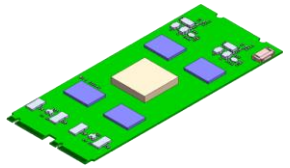

OCP modular form factors

Attendees and Participants

Semi Vendors
IP providers, EDA
Service providers

Tools, Manufacture,
Design, Test,
Integration

Systems vendors,
End users, ISVs,
Service Providers

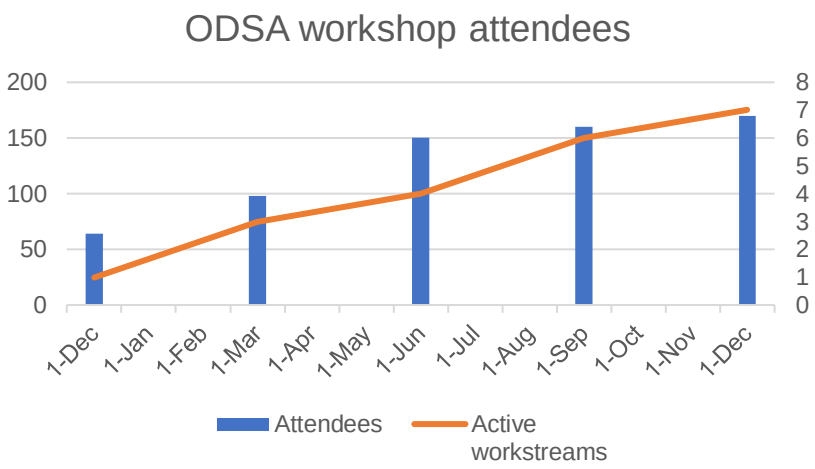

Attendance and/or participation do not imply corporate endorsement

A growing community

ODSA Work Streams

ODSA meets weekly Fridays, 8 AM Pacific
Each group meets weekly, details at https://www.opencompute.org/wiki/Server/ODSA

Workstream	Leader	Participants	Objective
PHY Layer	Robert Wang	IBM, Xilinx, Kandou Bus, Synopsys, Marvell, Google	PCIe PIPE adapter
Bunch of Wires	Mark Kuemerle	Marvell, Xilinx, Synopsys, Keysight Technologies, Google, IBM, SiFive	Scalable low-cost D2D PHY
CDX	Jawad Nasrullah	zGlue, ANSYS, Advanced Semiconductor Engineering, Inc., AyarLabs, cadence, Thrace Systems, Synopsys	Chiplet design exchange
Business	Sam Fuller	NXP, facebook, Microsoft Azure, AyarLabs	Chiplet workflowP
PoC hardware	JP Balachandran	cisco, facebook, Achronix Data Acceleration, Lattice, NXP	PoC board design
PoC software	Kevin Drucker	facebook, MACOM, Netronome, zGlue, socionext	Application/Infra software
Link layer	Open	Xilinx, NXP, intel, Ventana Micro, Microsoft Azure, cadence	ODSA Stack
OpenHBI	Kenneth Ma	Xilinx, cadence, IBM, Keysight Technologies, Synopsys, Samsung Semiconductors	High perf D2D PHY

Sub-project lead: Bapi Vinnakota BROADCOM

Attendance and/or participation do not imply corporate endorsement

ODSA Open D2D Interface Principles

- Layered design: Each layer evolves independently. Aim for at least one fully-open protocol option in each layer.
- Provide multiple options: Low-cost and complex packaging. Enable trade-offs for design cost and software impact. Market chooses best fit.
- Reuse: Leverage existing interfaces for abstraction between layers. Define interface adapters from existing interfaces to open PHYs (AIB, BoW, OpenHBI, XSR)
- Easy adoption: Carry popular protocols over D2D interfaces. PCIe/CXL for board aggregation and AXI/TileLink for SoC disaggregation. Define protocol adapters and enable reuse of off-the-shelf IP

ODSA D2D Interface

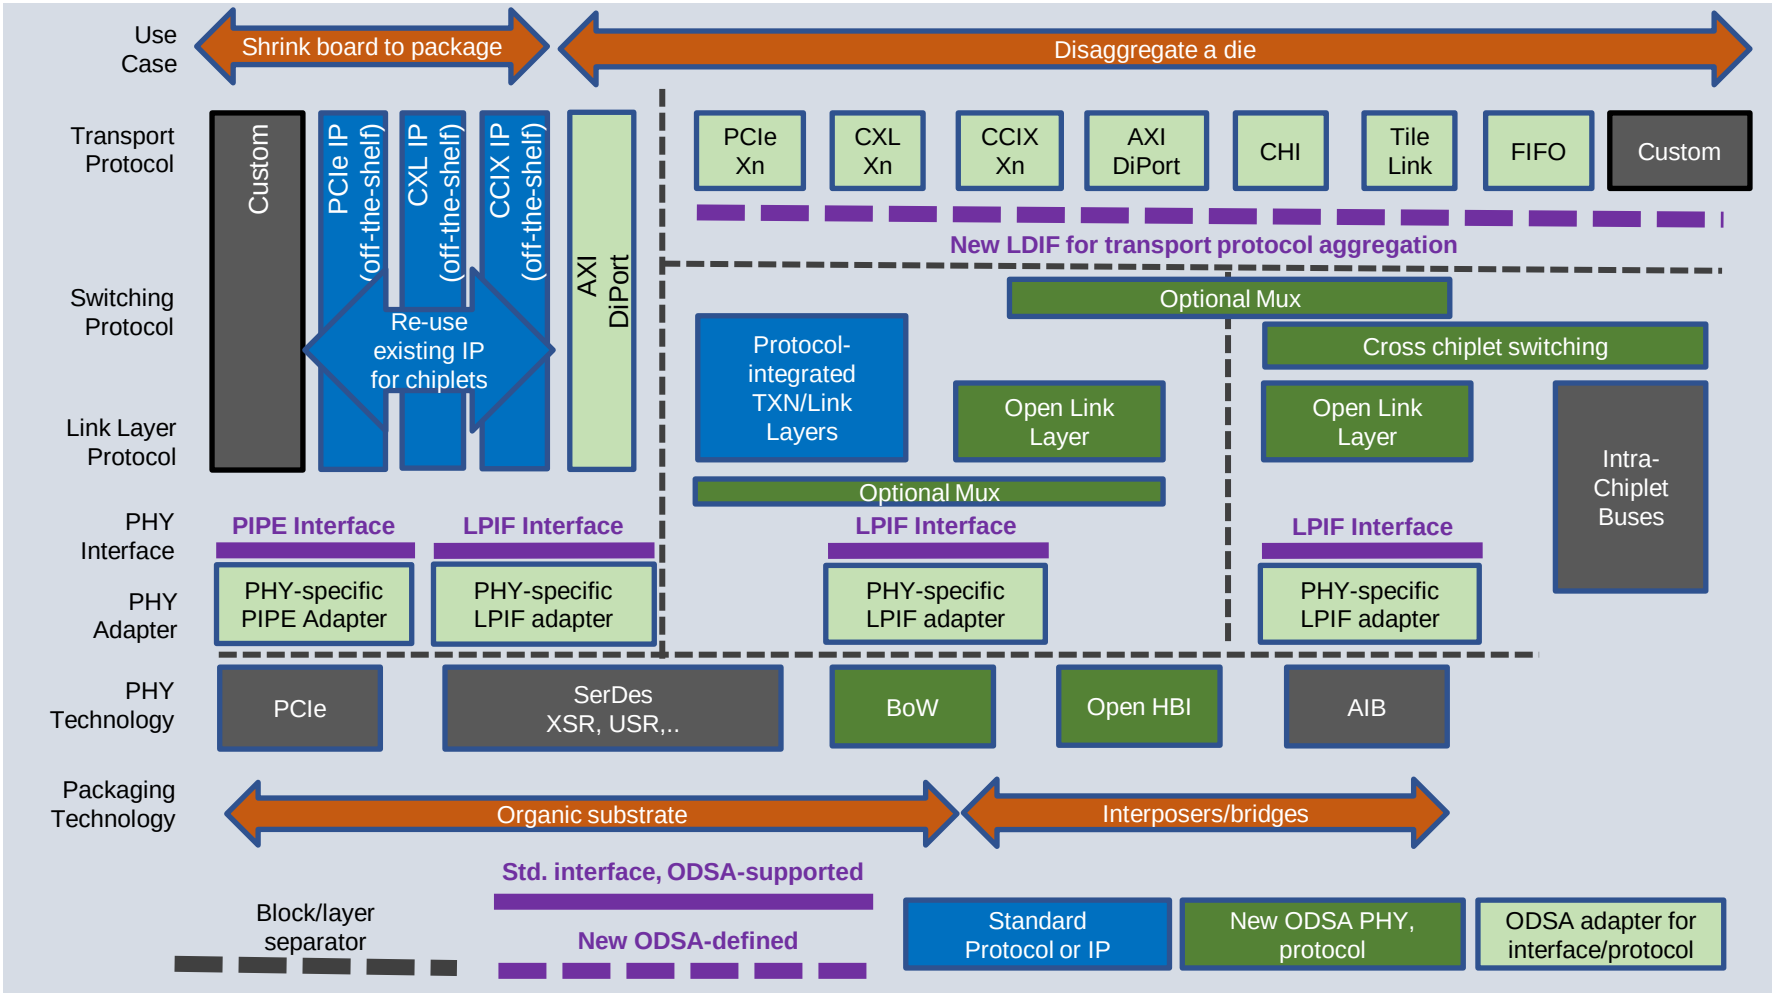

PHY: Bunch of Wires

Only open D2D PHY to offer graceful tradeoff of performance for complexity

Simple base design – 4-8 Gbps/wire

- Supports process nodes from 3nm to 65nm to enable heterogeneous designs.
- Beachfront bandwidth of 250 Gbps/mm of die edge with regular bumps at 130u spacing and simple laminate packaging
- 0.5 pJ/bit, 10x improvement over off-package PHYs

Increase beachfront density with [1,2] :

- More complex design – fast mode doubles bandwidth per lane
- More complex packaging – use microbumps at 50u spacing and wafer-level fanout/interposers
- Increase beachfront bandwidth to up to 2 Tbps/mm

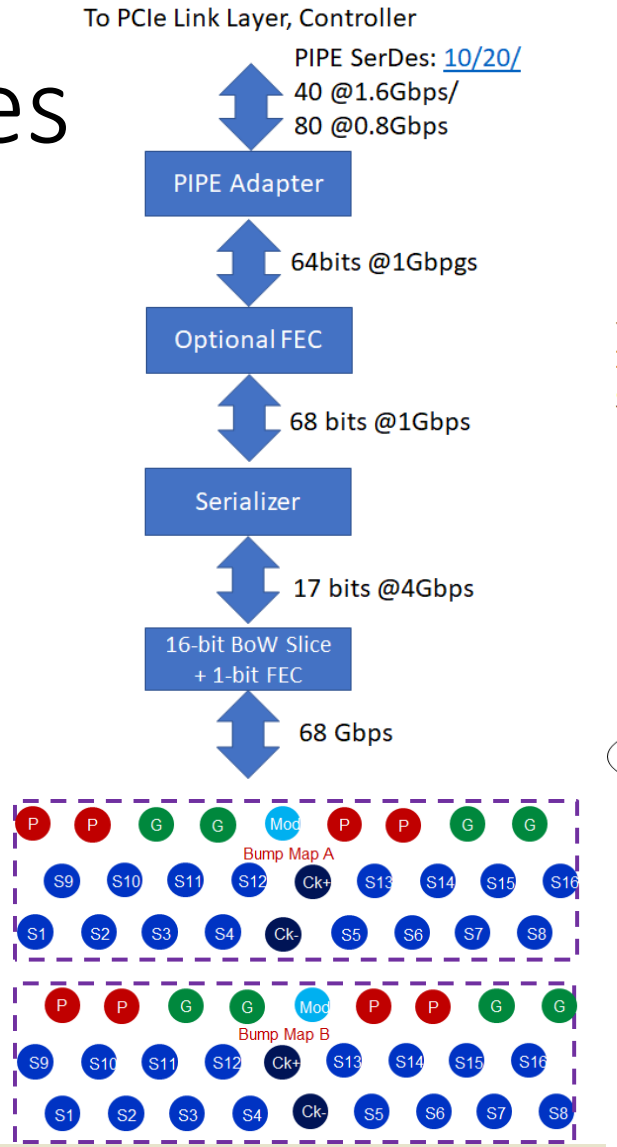

Base unit: 16-bit slice
~4xPCIe Gen 4 lanes

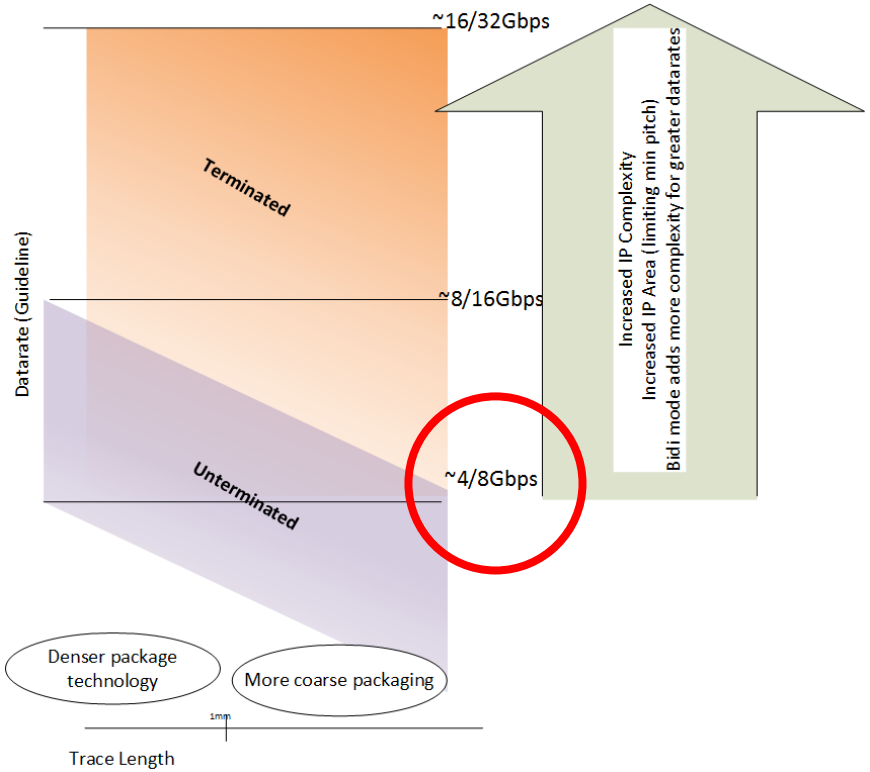

1. IEEE Micro, Jan. 2020,
2. https://github.com/opencomputeproject/ODSA-BoW

Logical Transactions: PIPE PCIe, CXL/DiPort AXI

PCIe, CXL over D2D interface

D2D PHY mimics a PCIe PHY to the controller

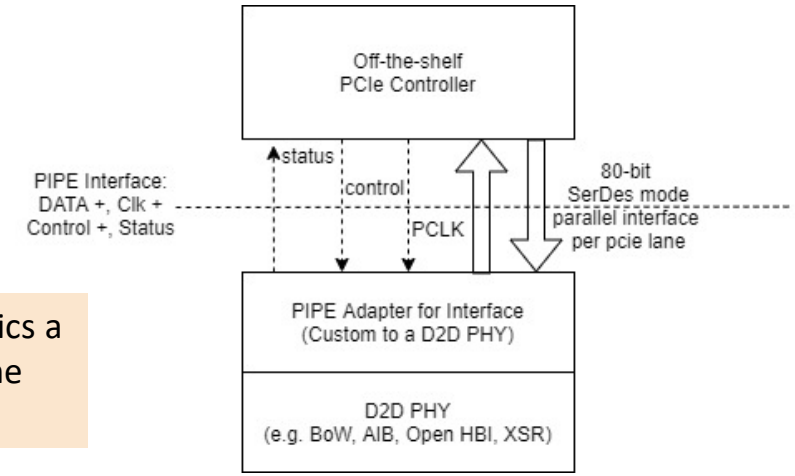

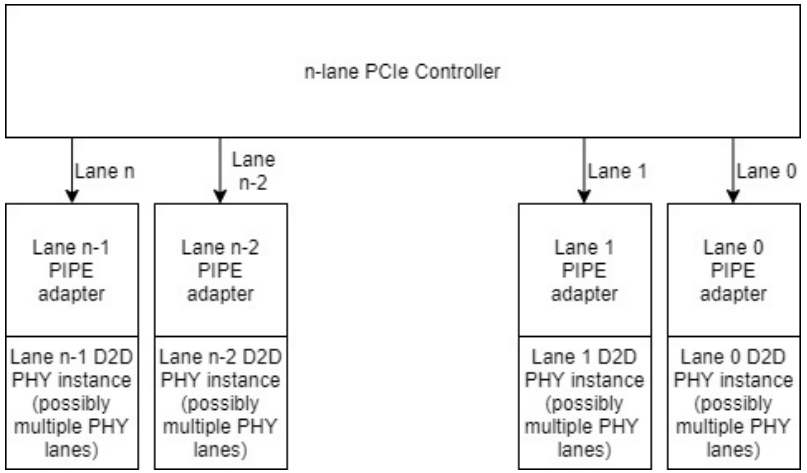

DiPort: AXI over D2D interface

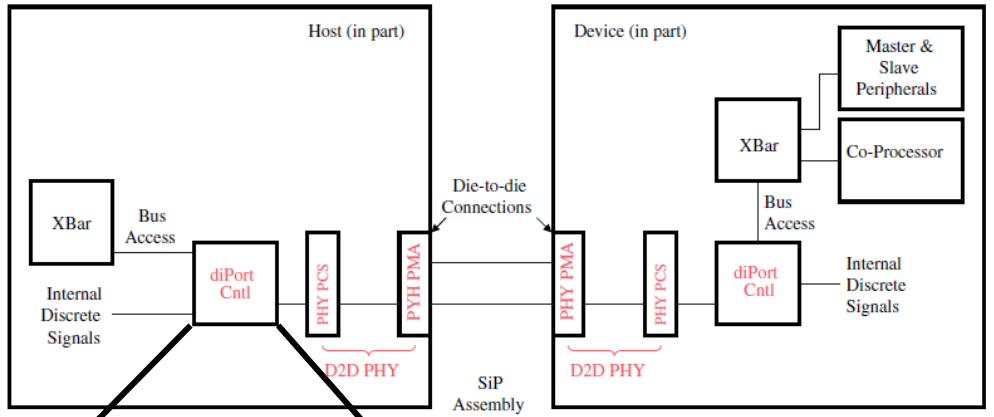

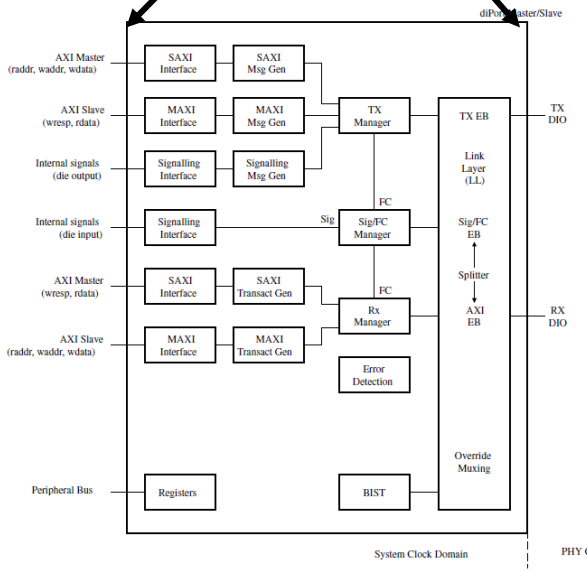

DiPort controller packetizes data and control signals

2020: ODSA PHY/Logic D2D interface

Complete open ODSA logic/PHY interfaces for chiplets for low-cost package.

- PCIe/CXL over BoW through standard PIPE interface
- AXI over BoW

<u>Combine the most common system (PCIe/CXL) or SoC protocols (AXI) with new ODSA protocols for chiplets.</u>

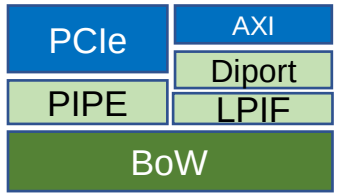

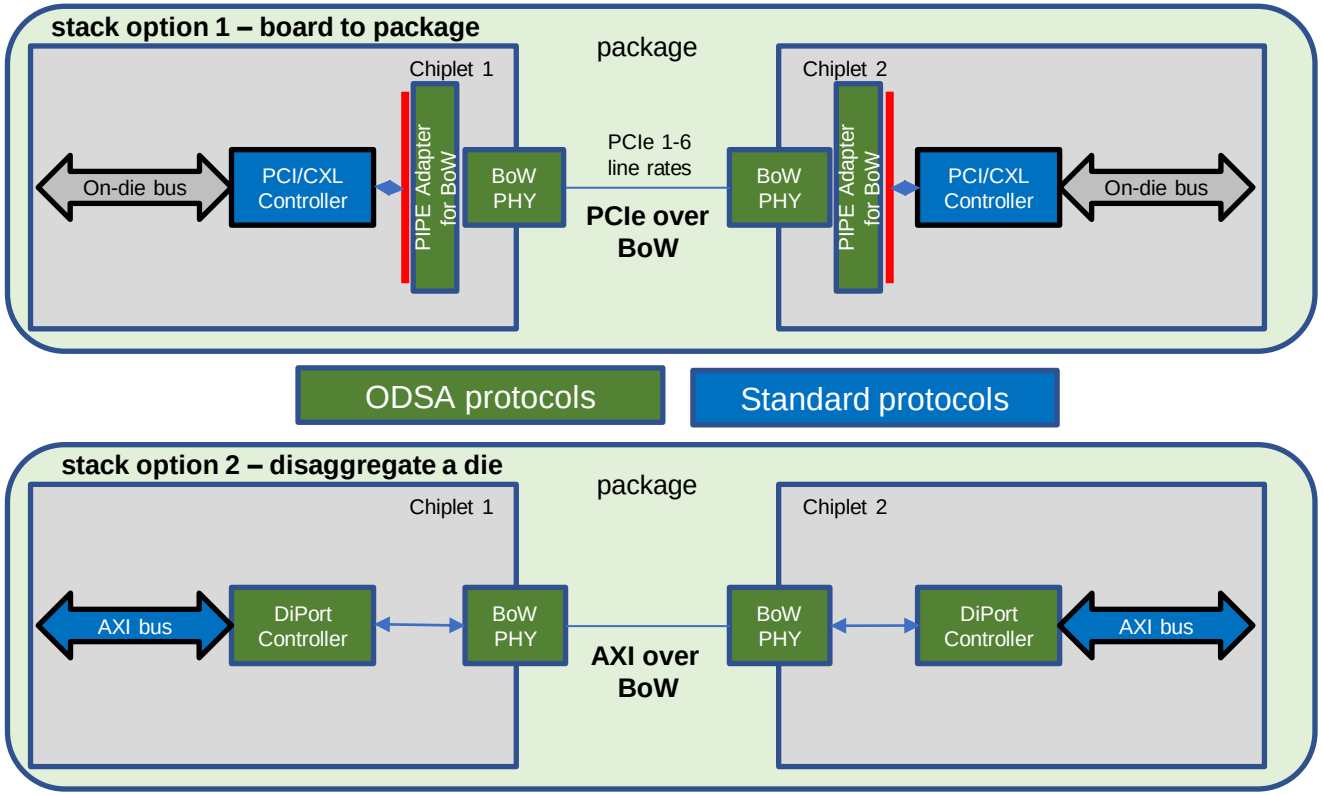

Proofs of Concept (PoCs) Development Plan

Workflow PoC to explore accelerator design[1]

Workflow PoC: Now

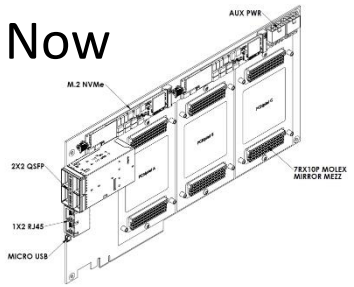

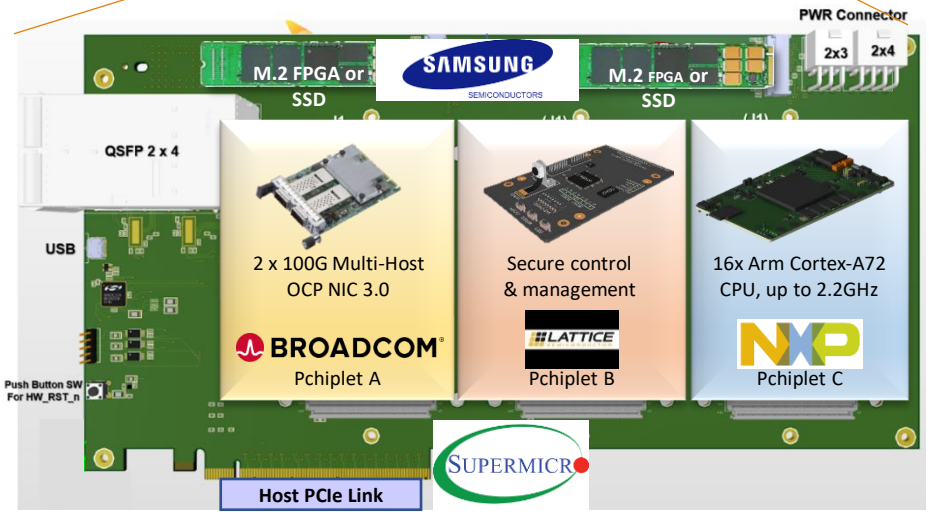

Chips to Chiplets Open Workflow

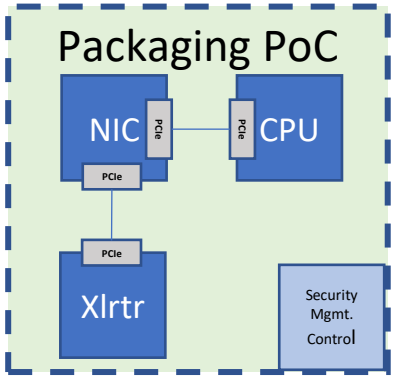

Workflow PoC Infrastructure, DSA Software Q4, 2020

D2D optimized Chiplets Workflow

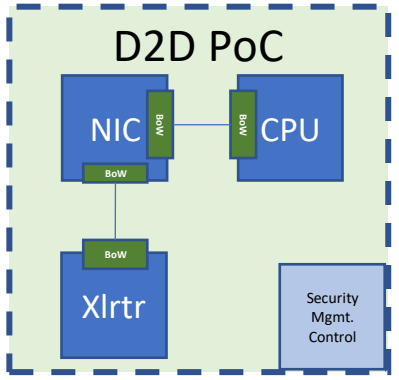

D2D test chips 2021

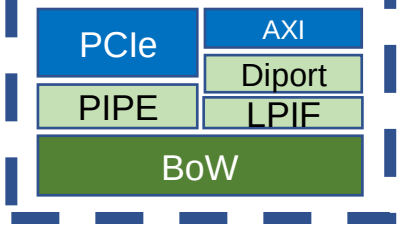

Current ODSA activities

Roadmap

1. OCP Virtual Summit, Keynote Session

Please Join Us

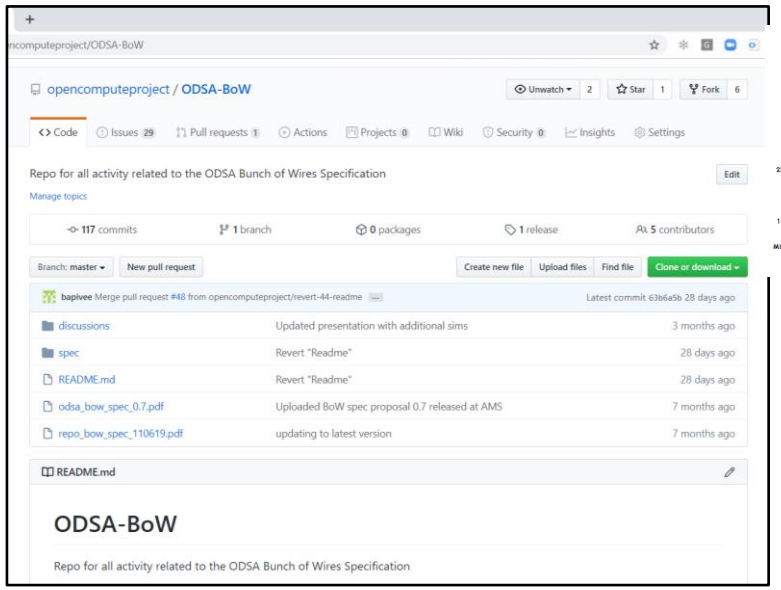

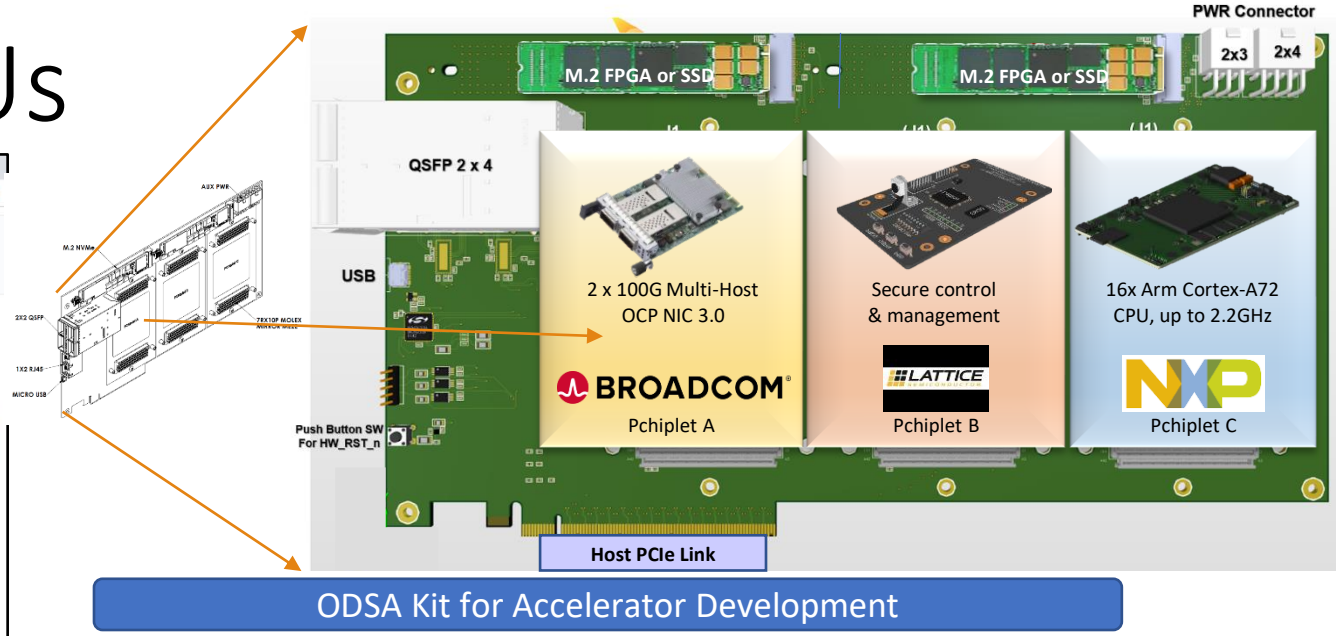

ODSA Kit for Accelerator Development

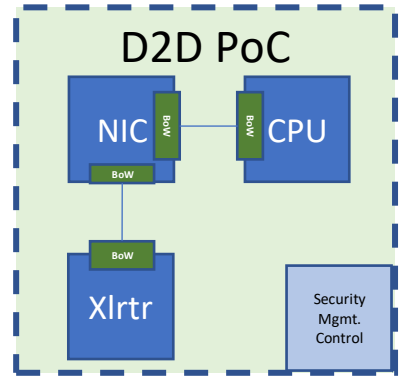

ODSA D2D PIPE Interface Specification

Chiplet-based designs have recently received a substantial amount of attention. In a chiplet-based product, a design is realized across multiple die integrated into one package. Chiplets require common physical and logical inter-chiplet interfaces to be integrated into one package. Chiplet-based designs use die2die (D2D) PHYs optimized for on-package data transport. D2D PHYs differ from those of both on-die highly parallel buses and off-package SerDes PHYs. Chiplet-based designs also require logical data protocols between the chiplets in a package. Current chiplet-based designs largely use closed physical and logical interfaces between the chiplets in a package and usually integrate internally-developed chiplets. With open interfaces, chiplets from multiple vendors can be integrated into one package.

The PIPE Interface

D2D technology is still evolving with no one interface being dominant. For example, the Open Domain-Specific Architecture (ODSA) project recently introduced two new open D2D PHYs. An abstraction layer above the PHY will help PHY and logical protocols for the chiplet interface evolve independently. The PIPE interface (latest version 5.2 as of November, 2019) is an abstraction layer between a PCIe PHY and the PCIe link layer developed by Intel. The interface maps a lane of PCIe traffic onto a PCIe PHY. The PIPE interface helps verify interoperability between PCIe controllers and PCIe PHYs (each usually made by different vendors).

This document specifies the PIPE interface for use as an abstraction layer for D2D interfaces in chiplet-based products. The first version of this document specifies how transaction protocols

OPEN Compute Project

DIPORT CONTROLLER DESIGN GUIDE

DRAFT

Many options to work together

- Join a spec effort
- Build an accelerator
- Prototype a package
- Develop a test chip

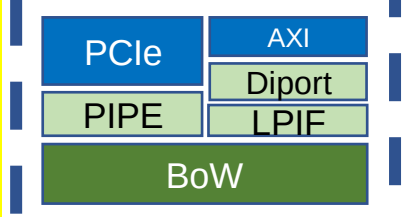

For More Information

- ODSA Wiki: All workshops, weekly calls (open access) with talks from industry experts.
- Specification proposals – workstreams meet weekly
  - Bunch of Wires GitHub repo (open access)
  - PIPE adapter, DiPort (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
  - ODSA white paper , ODSA Wiki
  - R. Farjadrad, M. Kuemerle, B. Vinnakota, "A Bunch-of-Wires (BoW) Interface for Interchiplet Communication", IEEE Micro, Jan. 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, R. Sethuram, B. Vinnakota, "Test Challenges in a Chiplet Marketplace", VLSI Test Symposium, Apr. 2020
  - B. Vinnakota, "The Open Domain-Specific Architecture: Year in Review", OCP Virtual Summit
  - BoW update and ODSA Overview at Hot Interconnect 2020

An Open Chiplet Marketplace

- First open physical and logical interface for chiplets
  - Uses abstraction interfaces to allow each layer to evolve independently
- Reusing current open transaction protocols support modular design
  - Supports a gradual migration to chiplets, modular design with chiplets
- Two complete PHY/Logical transaction protocol stacks in 2020
  - BoW PHY enables the integration of heterogeneous die from multiple vendors