High-Bandwidth and Low-Latency Standardized Interconnect for an Open Chiplet Ecosystem

Dr. Debendra Das Sharma

Intel Senior Fellow and Chief Architect of I/O Technology and Standards

Intel Corporation

Agenda

- Interconnects in Compute Landscape
- On-Package Interconnects: Opportunities and Challenges
- Universal Chiplet Interconnect Express (UCIe): An Open Standard for Chiplets
- Future Directions and Conclusions

Taxonomy, characteristics, and trends of interconnects

Off-Package

On-Package

Category	Type and Scale	Data Rate/ Characteristics	PHY Latency (Tx + Rx)
Latency Tolerant (Narrow, very high speed)	Networking / Fabric for Data Center Scale	56/ 112 GT/s-> 224 GT/s (PAM4) 4-8 Lanes, cables/ backplane	20+ ns (+ >100 FEC)
Latency Sensitive (Wide, high speed)	Load-Store I/O Arch. Ordering (PCIe/ CXL / SMP cache coherency – PCIe PHY) Node (-> Rack)	32 GT/s (NRZ) -> PCIe Gen6 64 GT/s (PAM4) Hundreds of Lanes Power, Cost, Si-Area, Backwards Compatible, Latency, On-board -> cables/ backplanes	<10ns (Tx+ Rx: PHY-PIPE) 0-1ns FEC overhead
Latency Sensitive (super-wide, high speed)	Load-Store and proprietary	4 G – 32G (single-ended, NRZ) 2D, 2.5D (-> 3D) Thousands of Lanes Ultra low power, ultra low latency High b/w density	<2ns (PHY – Transaction Layer)

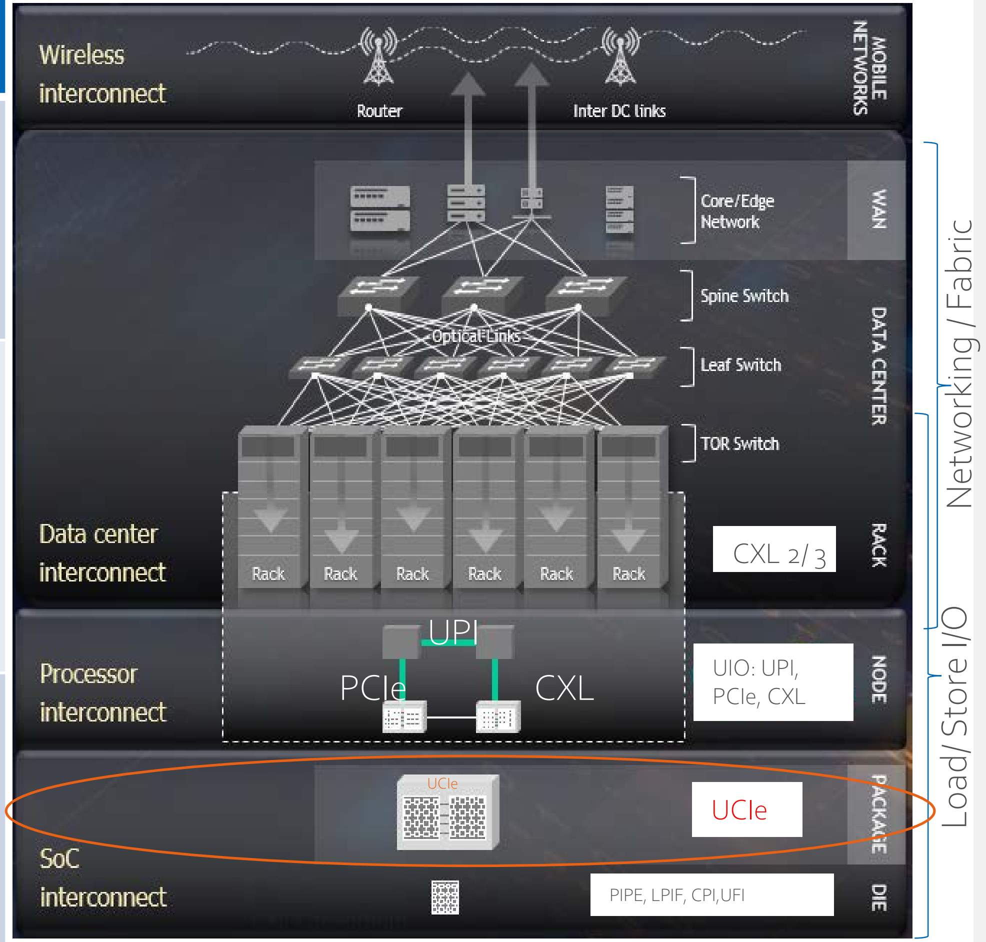

Load-Store I/O: From Package/ Node to Rack / Pod

Load-Store Interconnects : PCIe and CX

With PCIe: (900+ member companies)

- Memory Connected to CPU – Cacheable
- Memory Connected to PCIe device is Uncacheable
- Different Ordering rules across I/O vs coherency domains
- Ubiquitous I/O for compute continuum

With CXL: (~200 member companies)

- Caching and memory protocols on top of PCIe
- Device can cache memory
- Memory attached to device is cacheable
- Leverages PCIe infrastructure

PCIe and CXL very successful industry standards:

- Multi-generational, backward compatible, IP/ tools
- Compliance program with plug-and-play

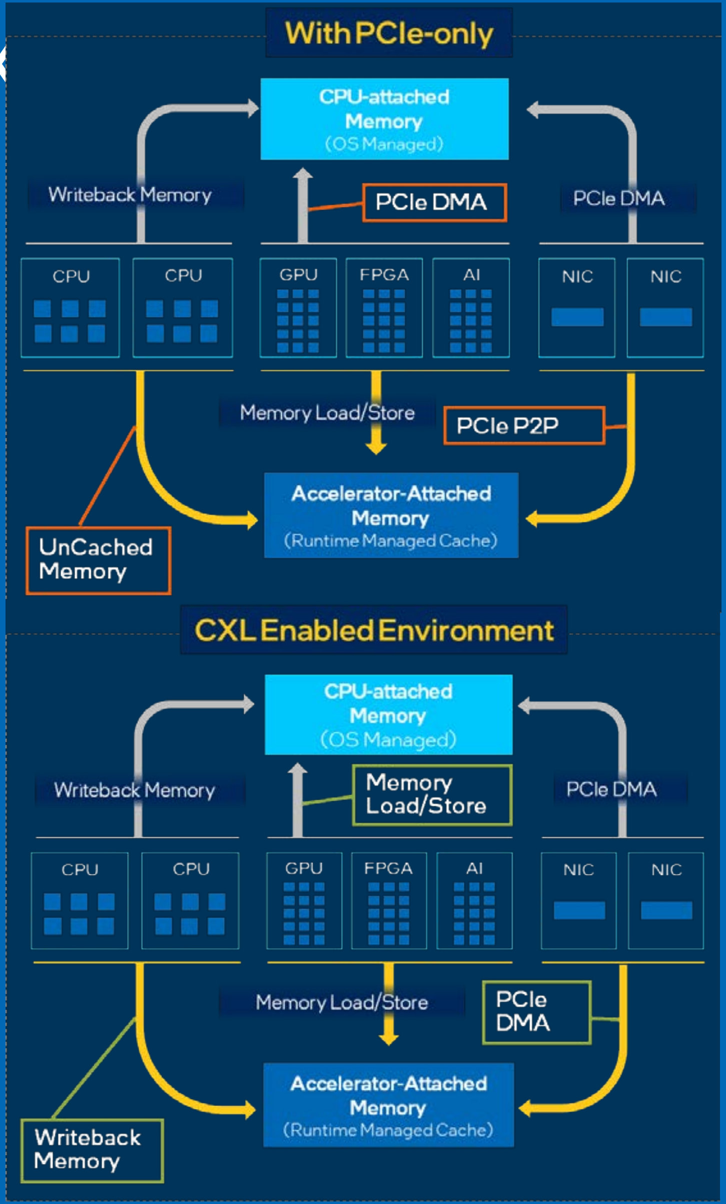

On-Package Interconnects should leverage PCIe/CXL infrastructure for standardization and Load-Store Usages.. Need to seamlessly move functionality from node to package to die level

Design Choice: Seamless Integration from Node → Package → On-die enables Reuse, Better User Experience

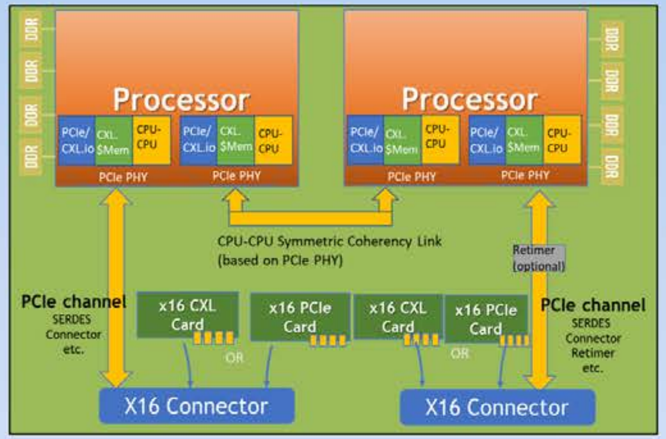

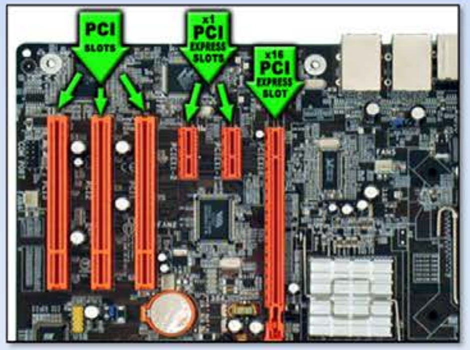

Node / Board Level Integration

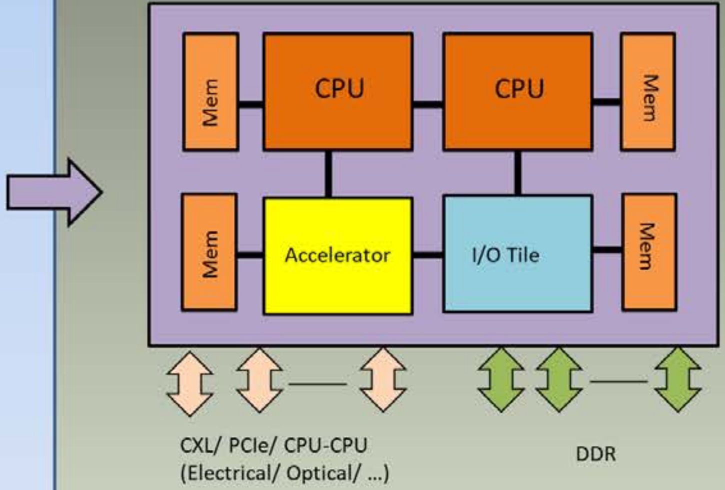

Package Level Integration
(with on-package interconnects)

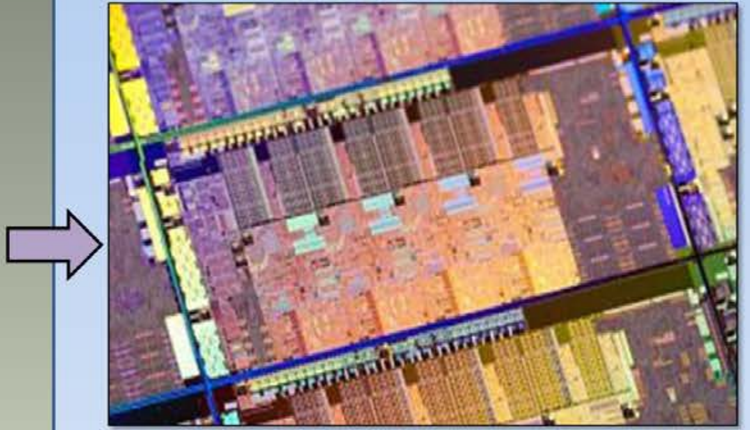

On-die Integration

Same Software, IP, and Subsystem to build scalable solutions offers economies of scale , time to market advantage, and seamless user experience. Innovations at the open slot in board level needs to migrate to package level for multiple usages!

Agenda

- Interconnects in Compute Landscape
- On-Package Interconnects: Opportunities and Challenges
- Universal Chiplet Interconnect Express (UCIe): An Open Standard for Chiplets
- Future Directions and Conclusions

Moore Predicted "Day of Reckoning"

"It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected[1]."

-Gordon E. Moore

[1]: "Cramming more components onto integrated circuits", Electronics, Volume 38, Number 8, April 19, 1965

Drivers for On-Package Chiplets

- Reticle Limit, yield optimization, scalable performance
    - Same dies on package (Scale-up)
- Increasing design costs at leading edge process nodes
    - Die-disaggregated dies across different nodes
    - Use new process node for advanced functionality
- Time to Market (Late binding)
- Custom silicon for different customers leveraging a base product
    - E.g., Different acceleration functions with common compute
- Different process nodes optimized for different functions
    - E.g., Memory, logic, analog, co-packaged optics
- High power-efficient bandwidth with low-latency access (e.g., HBM memory)

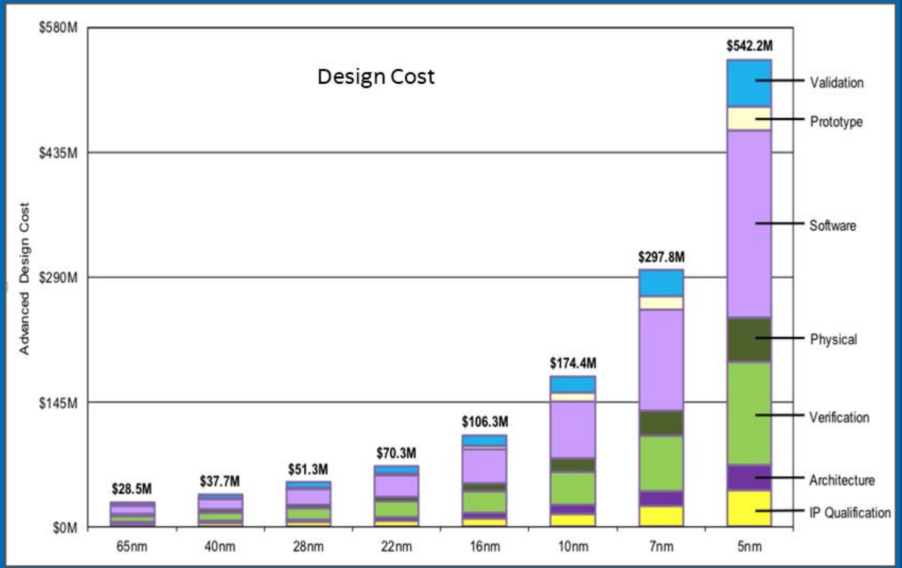

Source: IBS (as cited in IEEE Heterogeneous Integration Roadmap)

Components of Chiplet Interoperability

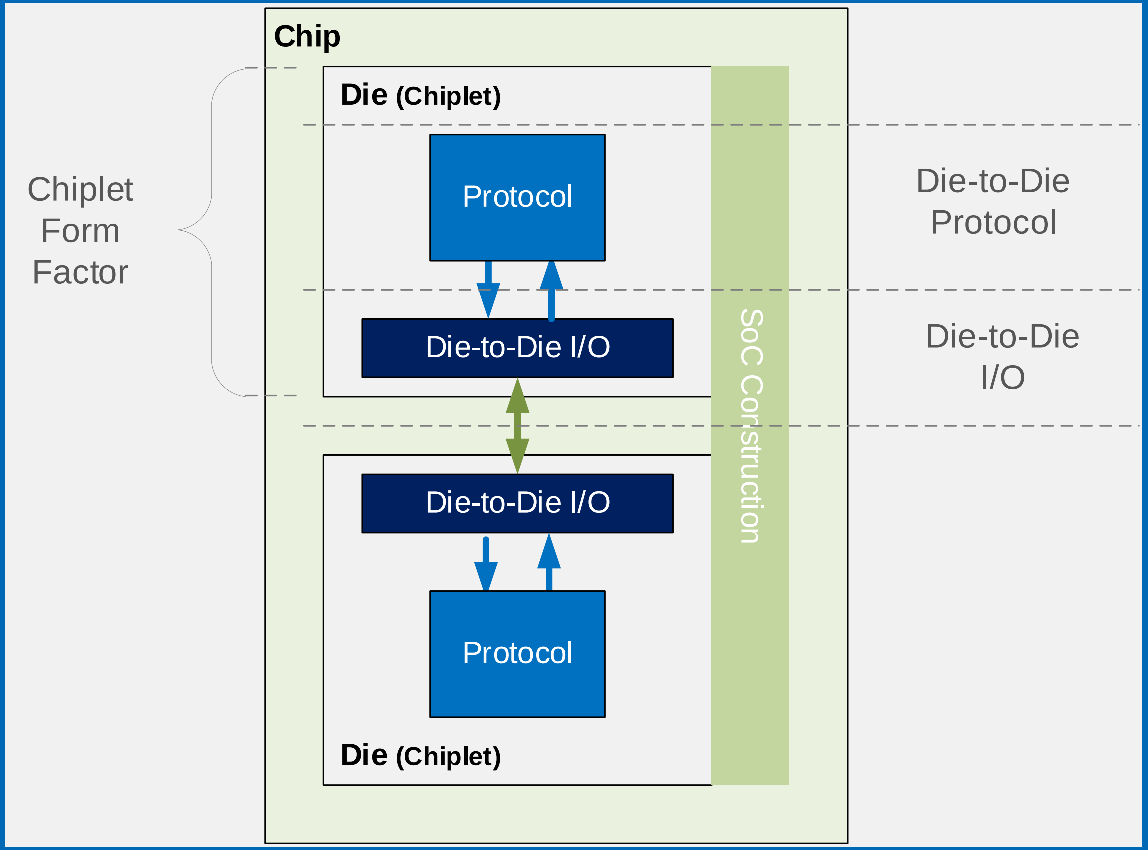

- Chiplet Form Factor
  - Die size
  - Bump location
  - Power delivery
  - Thermal characteristics
- SoC Construction (Application Layer)
  - SoC Reset
  - Initialization (e.g., fuses)
  - Register access
  - Security
- Die-to-Die Protocols (Data Link to Transaction Layer)
  - Link Layer, transaction Layer, etc.: PCIe/ CXL/ Raw/….
  - Internal Interface standardization for plug and play IPs
- Die-to-Die I/O (Physical Layer)
  - Bump arrangement and characteristics
  - Electrical & thermal characteristics
  - Substrate or interposer characteristics
  - Length budget, pJ/bit, bit error rate, …
  - Reset, clocking, initialization, and data transfer
  - Test and repair
  - Technology transition -> multiple bump arrangement/ frequency

Agenda

- Interconnects in Compute Landscape
- On-Package Interconnects: Opportunities and Challenges
- Universal Chiplet Interconnect Express (UCIe): An Open Standard for Chiplets
- Future Directions and Conclusions

Motivation for UCIe

OPEN CHIPLET: PLATFORM ON A PACKAGE

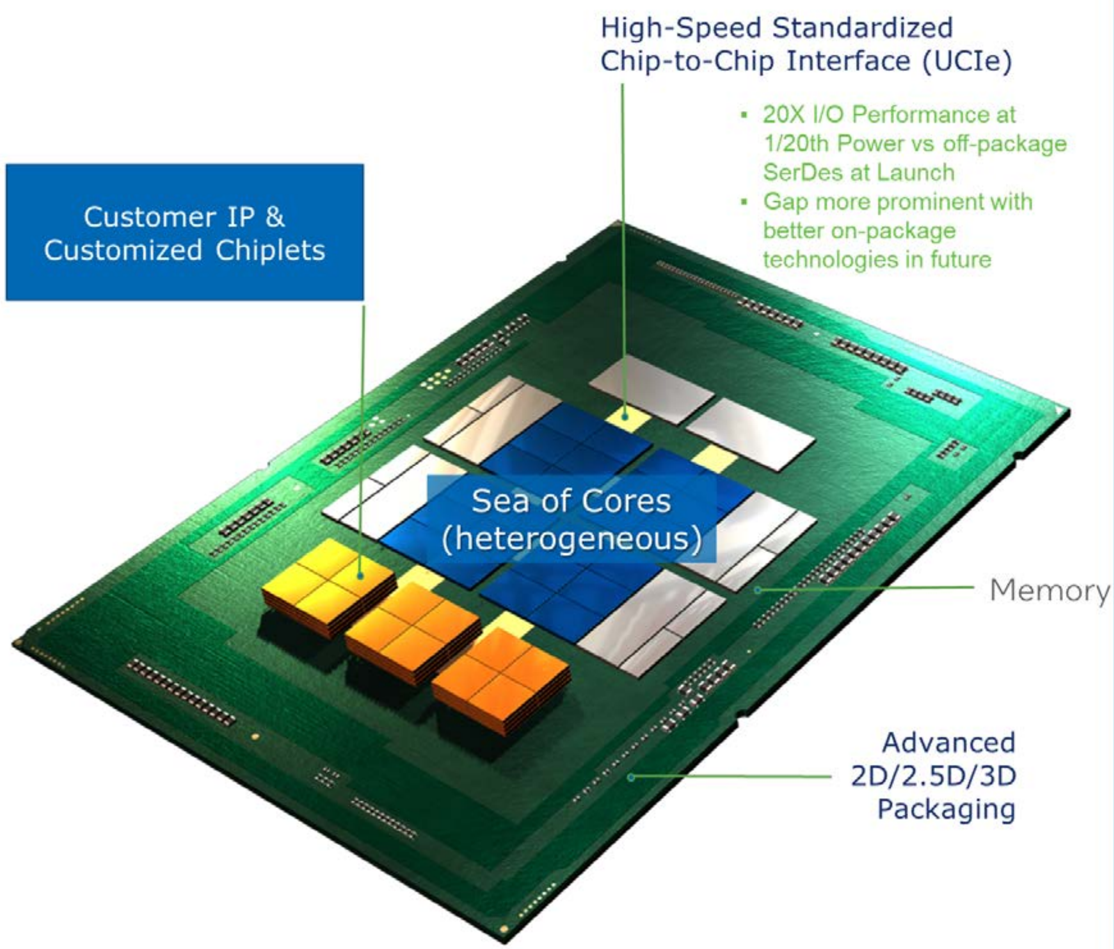

Heterogeneous Integration Fueled by an Open Chiplet Ecosystem
(Mix-and-match chiplets from different process nodes / fabs / companies / assembly)

Align Industry around an open platform to enable chiplet based solutions

- Enables SoC construction that exceeds maximum reticle size
  - Package becomes new System-on-a-Chip (SoC) with same dies (Scale Up)

- Reduces time-to-solution (e.g., enables die reuse)

- Lowers portfolio cost (product & project)
  - Enables optimal process technologies
  - Smaller (better yield)
  - Reduces IP porting costs
  - Lowers product SKU cost

- Enables a customizable, standard-based product for specific use cases (bespoke solutions)

- Scales innovation (manufacturing/ process locked IPs)

UCIe: Key Metrics and Adoption Criteria

Key Performance Indicators

- Bandwidth density (linear & area)
  - Data Rate & Bump Pitch
- Energy Efficiency (pJ/b)
  - Scalable energy consumption
  - Low idle power (entry/exit time)
- Latency (end-to-end: Tx+Rx)
- Channel Reach
  - Technology, frequency, & BER
- Reliability & Availability
- Cost: Standard vs advanced packaging

Factors Affecting Wide Adoption

- Interoperability
  - Full-stack, plug-and-play with existing s/w
  - Different usages/segments - ubiquity
- Technology
  - Across process nodes & packaging options
  - Power delivery & cooling
  - Repair strategy (failure/yield improvement)
  - Debug – controllability & observability
- Broad industry support / Open ecosystem
  - Learnings from other standards efforts

UCIe is architected and specified from the ground-up to deliver the best KPIs while meeting wide adoption criteria

UCIe 1.0 Specification

- **Layered Approach with industry-leading KPIs**
- **Physical Layer:** Die-to-Die I/O
- **Die to Die Adapter:** Reliable delivery
  - Support for multiple protocols: bypassed in raw mode
- **Protocol:** CXL/PCIe and Streaming
  - **CXL™/PCIe® for volume attach and plug-and-play**
    - SoC construction issues are addressed w/ CXL/PCIe
    - CXL/PCIe addresses common use cases
      - I/O attach, Memory, Accelerator
  - **Streaming for other protocols**
    - Scale-up (e.g., CPU/ GP-GPU/Switch from smaller dies)
    - Protocol can be anything (e.g., AXI/CHI/SFI/CPI/ etc)
- **Well defined specification**: interoperability and future evolution
  - Configuration register for discovery and run-time
    - control and status reporting in each layer
    - transparent to existing drivers
  - Form-factor and Management
  - Compliance for interoperability
  - Plug-and-play IPs with RDI/ FDI interface

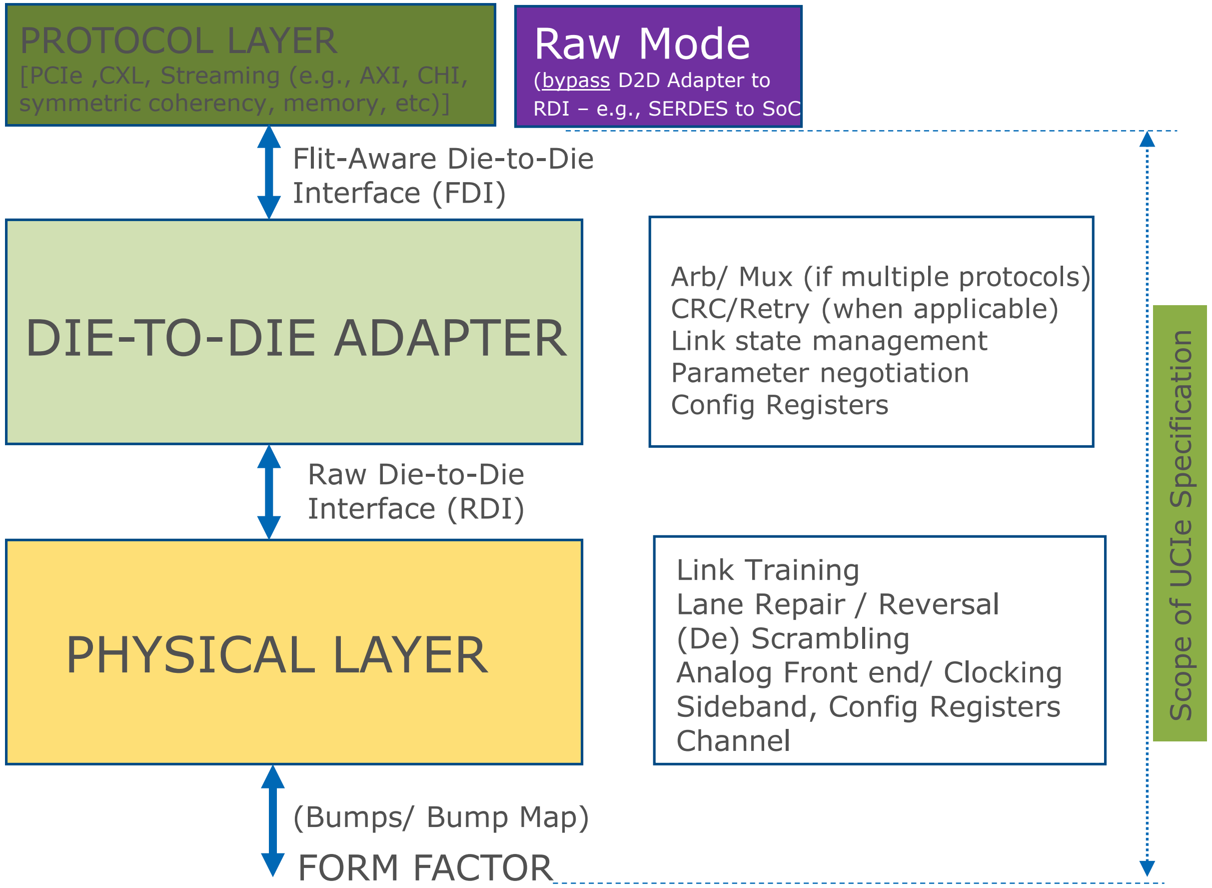

PROTOCOL LAYER
[PCIe ,CXL, Streaming (e.g., AXI, CHI, symmetric coherency, memory, etc)]

Raw Mode
(bypass D2D Adapter to RDI – e.g., SERDES to SoC

Flit-Aware Die-to-Die Interface (FDI)

DIE-TO-DIE ADAPTER

Arb/ Mux (if multiple protocols)
CRC/Retry (when applicable)
Link state management
Parameter negotiation
Config Registers

Raw Die-to-Die Interface (RDI)

PHYSICAL LAYER

Link Training
Lane Repair / Reversal
(De) Scrambling
Analog Front end/ Clocking
Sideband, Config Registers
Channel

(Bumps/ Bump Map)
FORM FACTOR

Scope of UCIe Specification

UCIe 1.0: Supports Standard and Advanced Packages

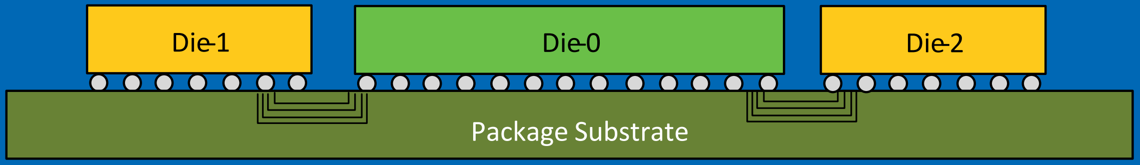

(Standard Package)

Standard Package: 2D – cost effective, longer distance

Advanced Package: 2.5D – power-efficient, high bandwidth density

Dies can be manufactured anywhere and assembled anywhere – can mix 2D and 2.5D in same package – Flexibility for SoC designer

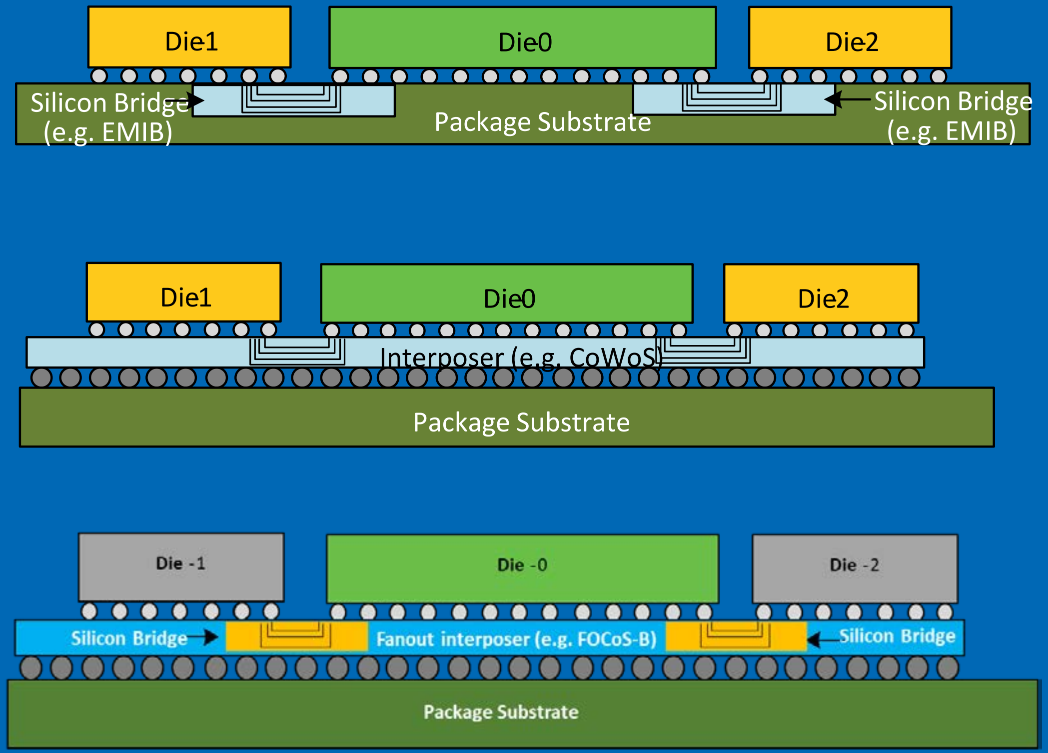

(Multiple Advanced Package Options)

One UCIe 1.0 Spec covers both type of packaging options

UCIe Usage Model: SoC at Package Level

- SoC as a Package level construct
  - Standard and/ or Advanced package
  - Homogeneous and/or heterogeneous chiplets
  - Mix and match chiplets from multiple suppliers
- Across segments: Hand-held, Client, Server, Workstation, Comms, HPC, etc
  - Similar to PCIe/ CXL at board level

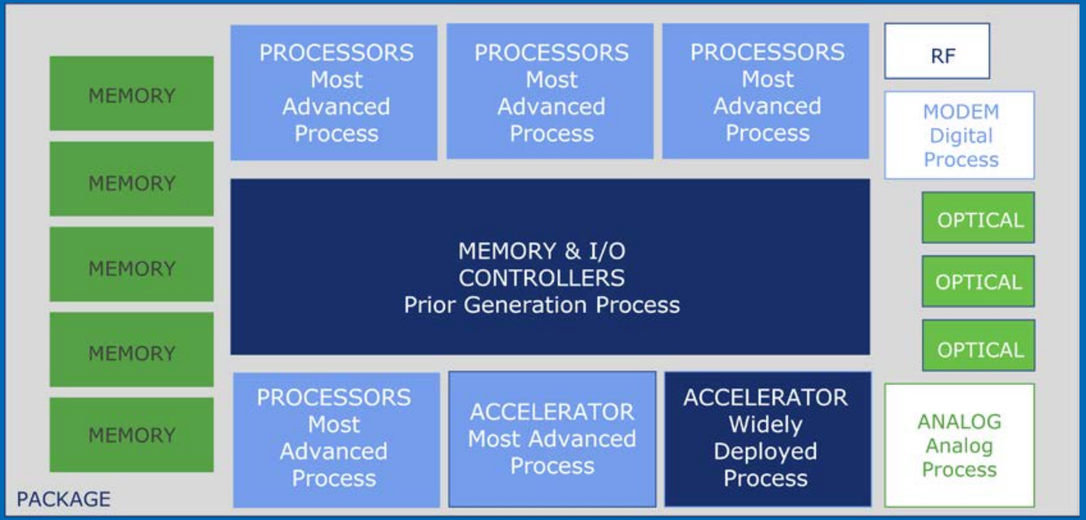
MEMORY
MEMORY
MEMORY
MEMORY
MEMORY
PROCESSORS Most Advanced Process
PROCESSORS Most Advanced Process
PROCESSORS Most Advanced Process
RF
MODEM Digital Process
MEMORY & I/O CONTROLLERS Prior Generation Process
OPTICAL
OPTICAL
OPTICAL
PROCESSORS Most Advanced Process
ACCELERATOR Most Advanced Process
ACCELERATOR Widely Deployed Process
ANALOG Analog Process
PACKAGE

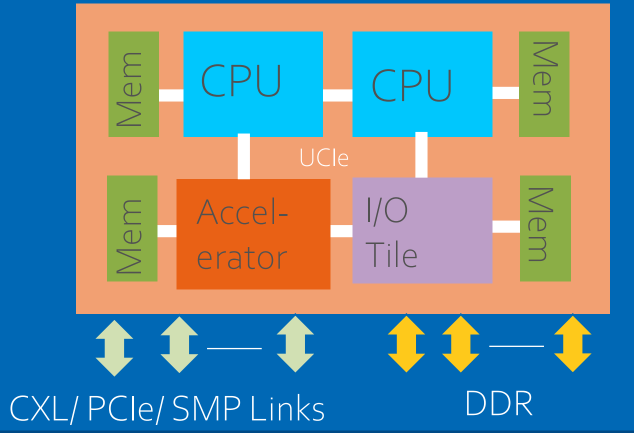
Mem
CPU
CPU
Mem
UCIe
Mem
Accel-erator
I/O Tile
Mem
CXL/ PCIe/ SMP Links
DDR

Example Scale-up SoC from homogeneous dies: Large Switch with on-die protocol as streaming over UCIe

- Need large radix CXL switches – challenges: reticle limit, cost, etc.
- UCIe based Chiplets should help with scalable products
- 64G Gen6 x16b CXL links
- UCIe as d2d interconnect – while this is a scale-up CXL switch , a switch vendor may prefer to have their on-die interconnect protocol be transported over UCIe rather than create a hierarchy of switches which will not work for CXL 2.0 tree-based topology

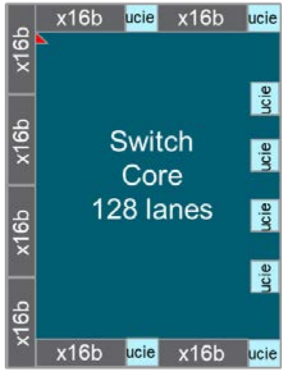

Small CXL Switch (128 lanes)

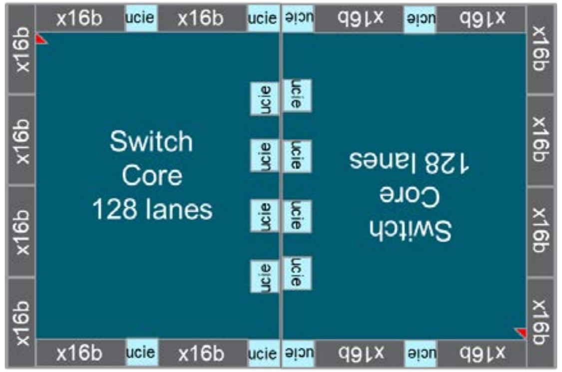

Medium-sized CXL Switch (256 lanes)

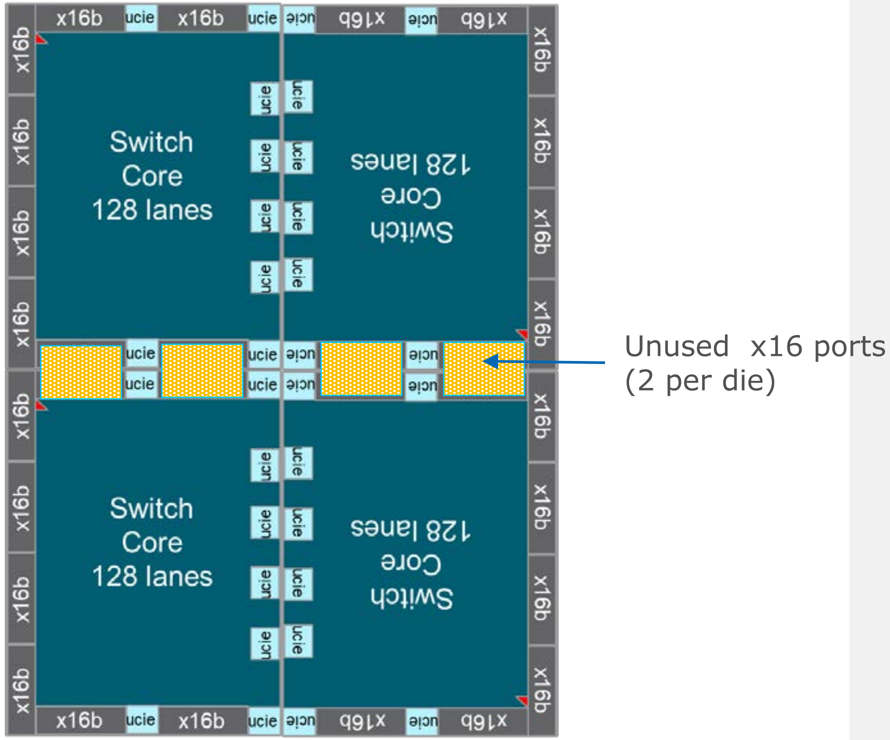

Large CXL switch (512 lanes)

One can construct CPUs (low, medium, large core-count CPUs) from smaller dies connected through UCIe using the same principle
Here the UCIe PHY and D2D adapter will carry the packetized version of internal CPU interconnect fabric

Ack: Nathan Kalyanasundaram

Example Scale-up Package using Streaming and open-plug-in using PCIe/ CXL

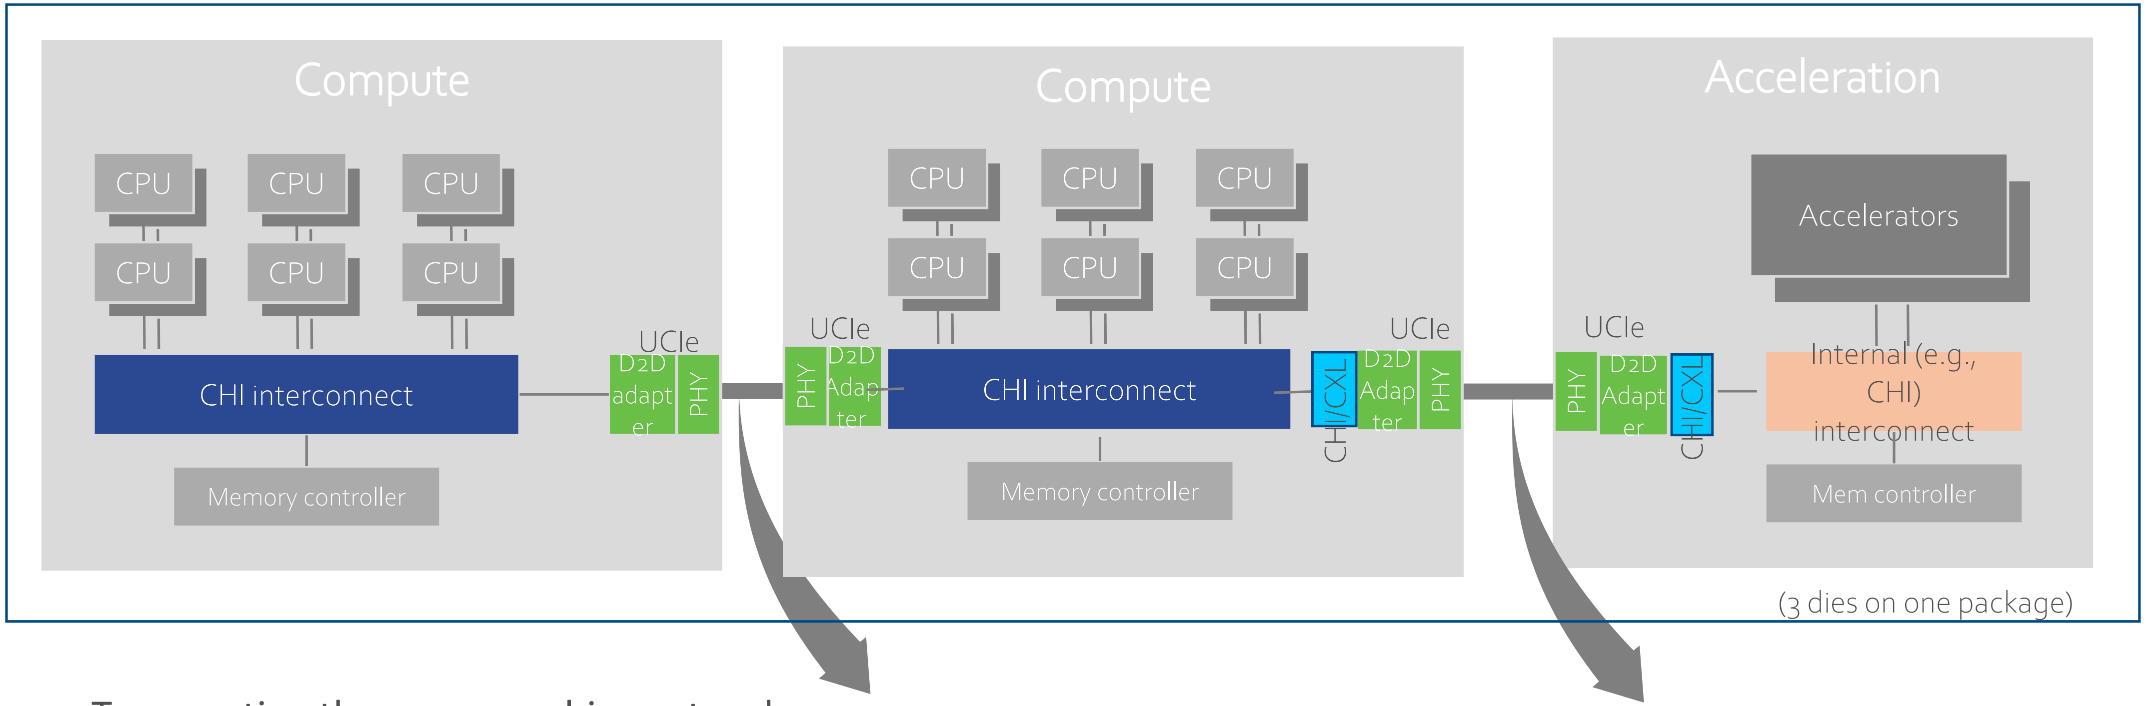

Not drawn to scale

- Transporting the same on-chip protocol allows seamless use of architecture specific features without protocol conversion
- Streaming interface with additional flit formats provide link robustness using UCIe defined data-link CRC and retry

- Any device type in this open plug-in slot with CXL (or CHI if both support it)

Ack: Marvin Denman, Bruce Mathewson, Francisco Socal, Durgesh Srivastava, Dong Wei

UCIe Usage: Off-Package Connectivity with UCIe Retimers

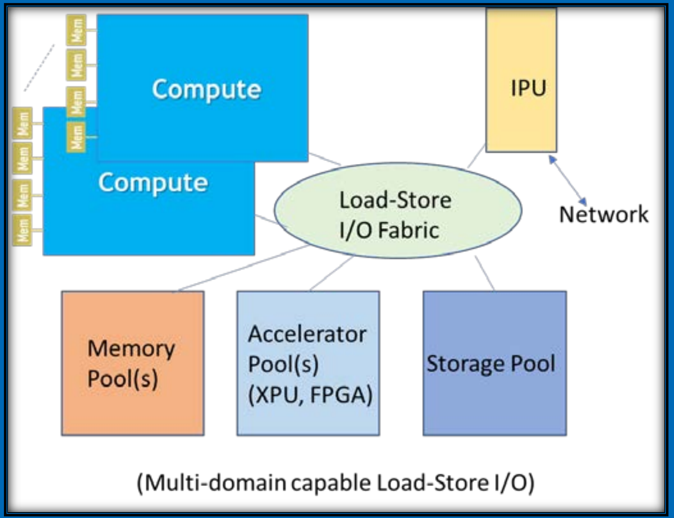

(Multi-domain capable Load-Store I/O)

(Vision: Load-Store I/O (CXL) as the fabric across the Pod providing low-latency and high bandwidth resource pooling/ sharing as well as message passing)

(Pod of Racks)

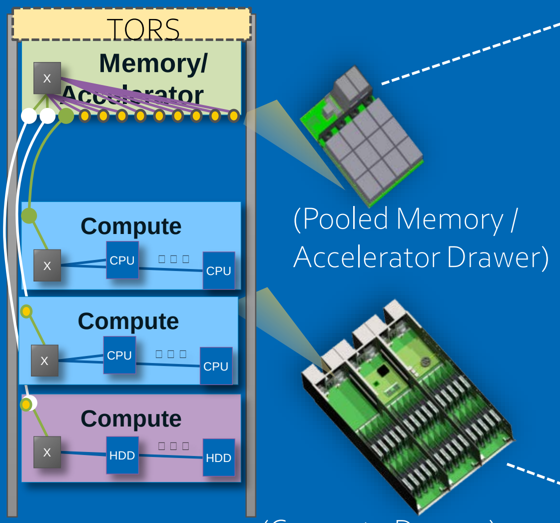

(Pooled Memory / Accelerator Drawer)

(Compute Drawer)

CXL Rack / Pod level connected using long-reach media (Electrical/ Optical/ ..) through UCIe Retimers (e.g., co-packaged optics)

(Physical Connectivity using UCIe-Retimer based co-packaged optics)

Rack/ Pod Level resource pooling/ sharing with UCIe

(Optical connections: Intra-Rack and Pod)

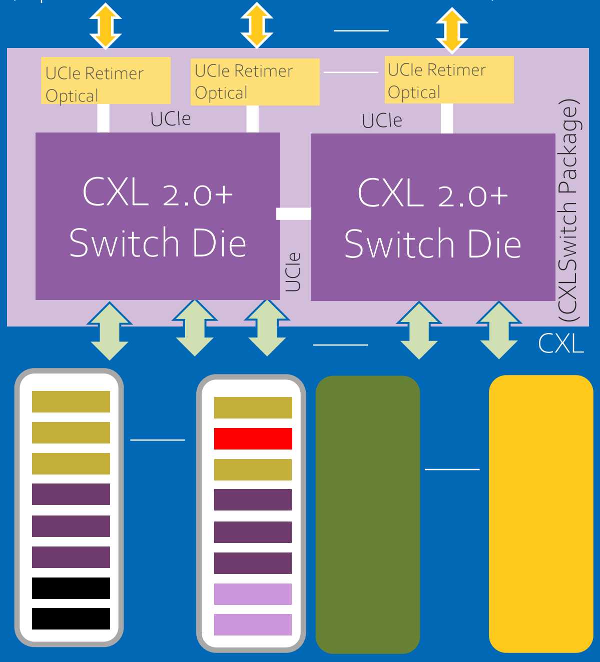

(Pooled/ Shared Memory)

(Pooled Accelerator)

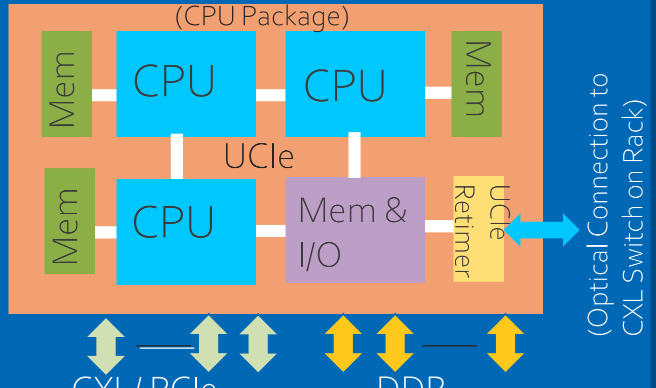

CXL/ PCIe

DDR

(Interconnects at drawer level)

UCIe 1.0: Characteristics and Key Metrics

CHARACTERISTICS	STANDARD PACKAGE	ADVANCED PACKAGE	COMMENTS
Data Rate (GT/s)	4, 8, 12, 16, 24, 32		Lower speeds must be supported -interop (e.g., 4, 8, 12 for 12G device)
Width (each cluster)	16	64	Width degradation in Standard, spare lanes in Advanced
Bump Pitch (um)	100 – 130	25 - 55	Interoperate across bump pitches in each package type across nodes
Channel Reach (mm)	<= 25	<=2	

KPIs / TARGET FOR KEY METRICS	STANDARD PACKAGE	ADVANCED PACKAGE	COMMENTS
B/W Shoreline (GB/s/mm)	28 – 224	165 – 1317	Conservatively estimated: AP: 45u; Standard: 110u; Proportionate to data rate (4G – 32G)
B/W Density (GB/s/mm^2)	22-125	188-1350	
Power Efficiency target (pJ/b)	0.5	0.25	
Low-power entry/exit latency	0.5ns <=16G, 0.5-1ns >=24G		Power savings estimated at >= 85%
Latency (Tx + Rx)	< 2ns		Includes D2D Adapter and PHY (FDI to bump and back)
Reliability (FIT)	0 < FIT (Failure In Time) << 1		FIT: #failures in a billion hours (expecting ~1E-10) w/ UCIe Flit Mode

UCIe 1.0 delivers the best KPIs while meeting the projected needs for the next 5-6 years.
Wide industry leader adoption spanning semiconductor, manufacturing, assembly, & cloud segments.

Ingredients of broad inter-operable chiplet ecosystem

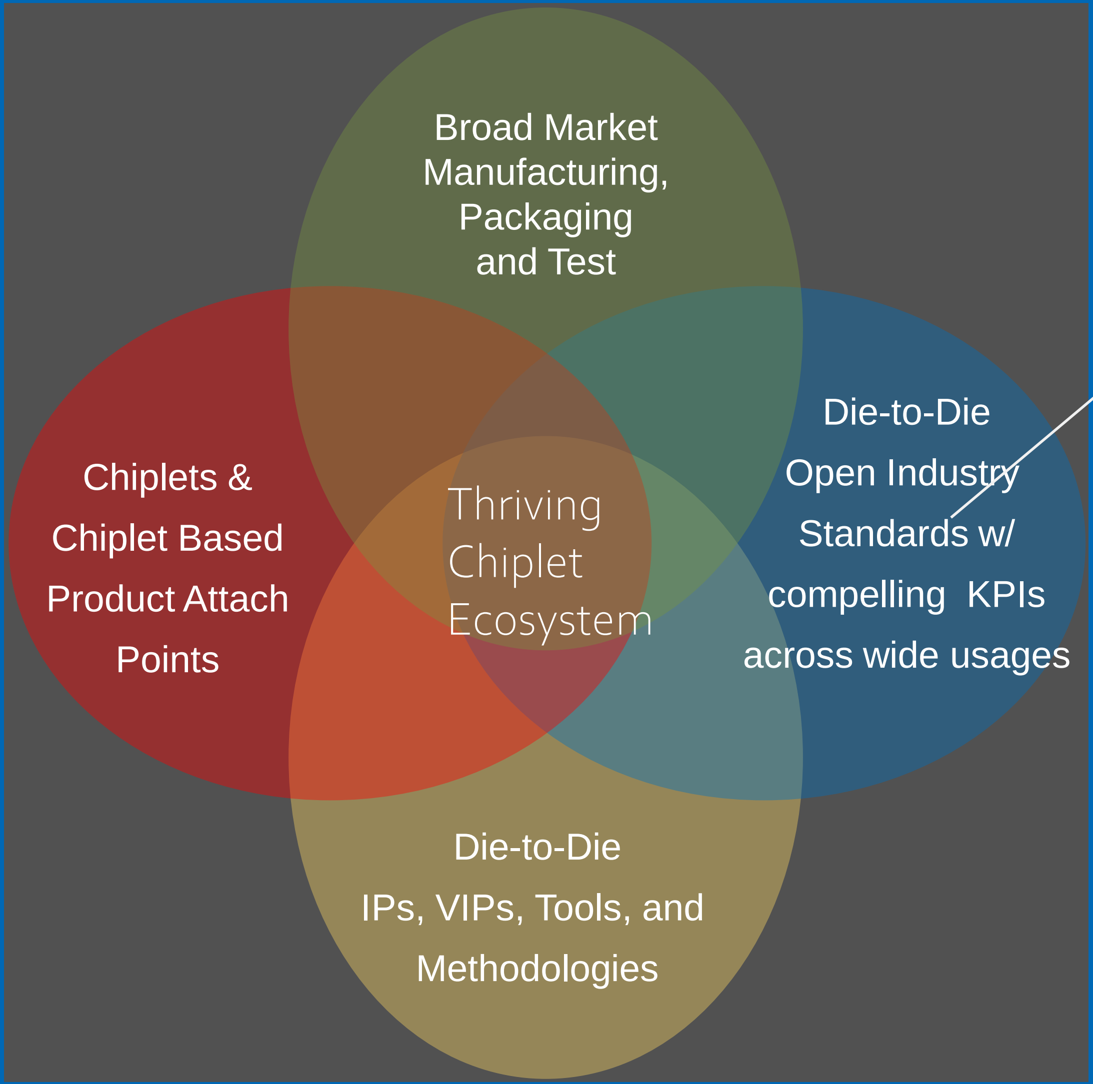

Broad Market Manufacturing, Packaging and Test

Chiplets & Chiplet Based Product Attach Points

Thriving Chiplet Ecosystem

Die-to-Die Open Industry Standards w/ compelling KPIs across wide usages

Die-to-Die IPs, VIPs, Tools, and Methodologies

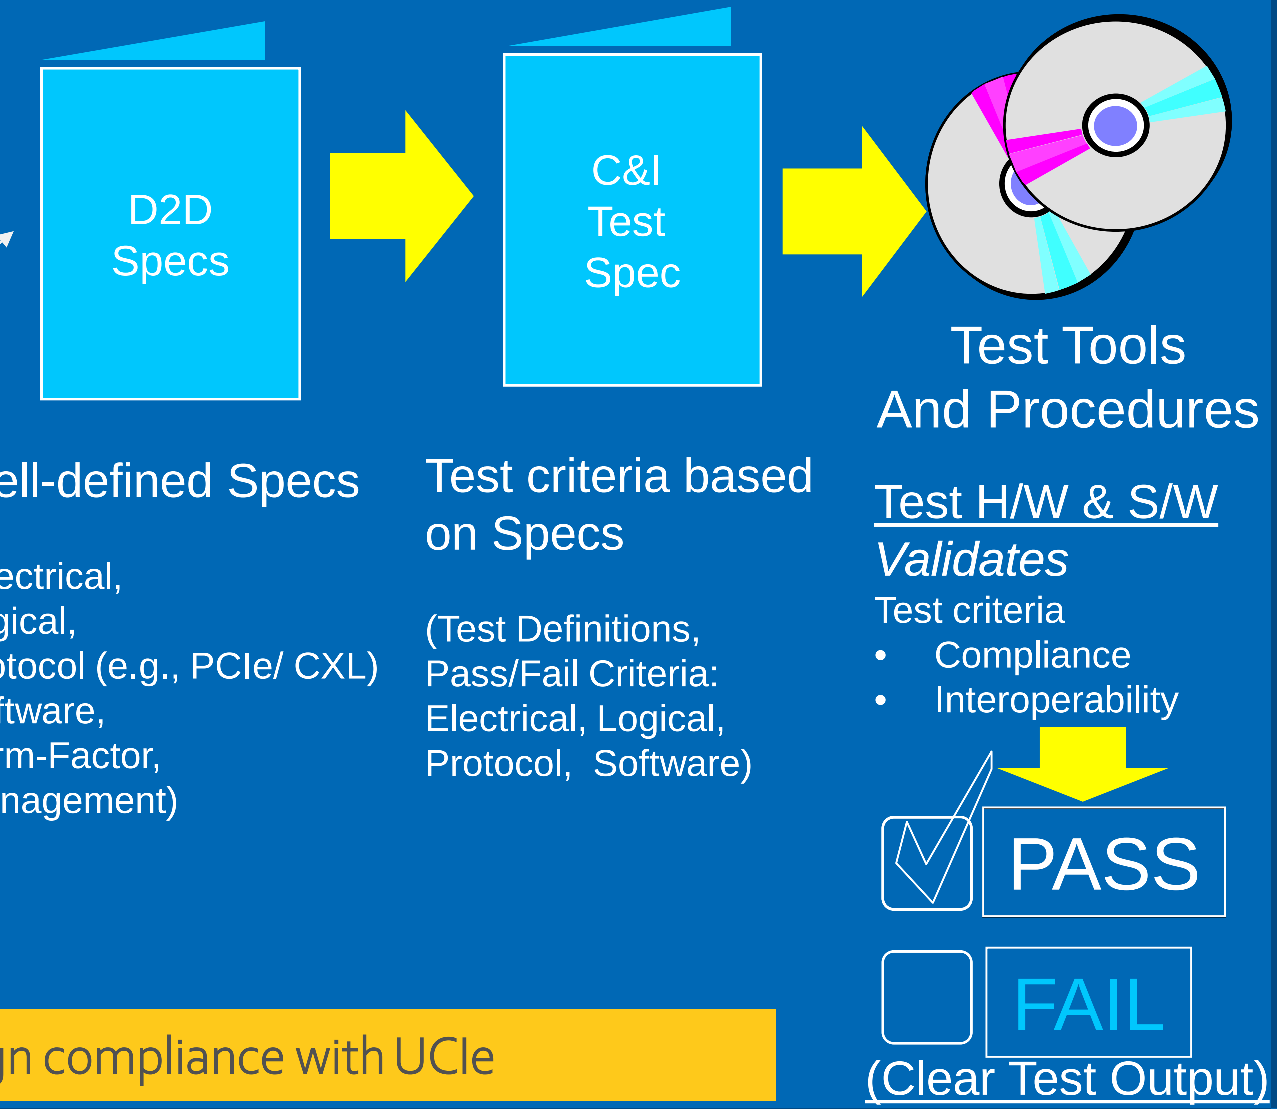

Well-defined Specs

(Electrical,
Logical,
Protocol (e.g., PCIe/ CXL)
Software,
Form-Factor,
Management)

Test criteria based on Specs

(Test Definitions,
Pass/Fail Criteria:
Electrical, Logical,
Protocol, Software)

Test H/W & S/W *Validates*
Test criteria
- Compliance
- Interoperability

☑ PASS

☐ FAIL

(Clear Test Output)

Predictable path to design compliance with UCIe

100+ Member Companies Board Members

Leaders in semiconductors, packaging, IP suppliers, foundries, and cloud service providers are driving the open chiplet ecosystem.

JOIN US!

Alibaba Group
AMD
arm
ASE GROUP
Google Cloud
intel
Meta
Microsoft
NVIDIA
Qualcomm
SAMSUNG
tsmc

Future Directions and Conclusions

Chiplets and D2D interface are essential to the compute continuum

- Power-efficient performance, yield optimization, different functions, custom solutions, cost-effective

UCIe standardization will propel the development an open ecosystem

- Open plug-and-play “slot” at package level will unleash innovations
- Evolution needs to track the underlying packaging technology to deliver compelling metrics
- Form-factor, New Protocols, and manageability are some other areas for innovation

The open chiplet journey with UCIe just started! Join us in what will be an exciting journey for decades!