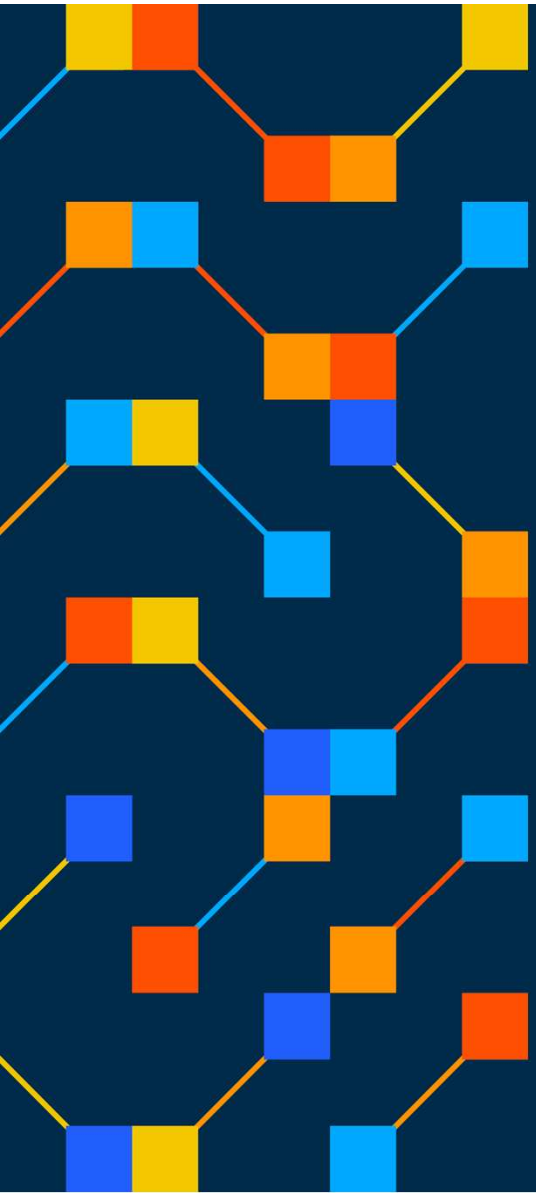

Multi-Die System IP

MpSoC 2025, Megeve France

Charlie Janac
President and CEO, Arteris

June, 17, 2025

ARTERIS IP

Today's Chips Have Become The Core Of Electronic Systems

System IP solutions address ~10-20% of silicon area*, solving IP reuse and SoC design challenges → Best Power/Performance/Area (PPA), lower costs, and reduced project risks

Requirements ▪ Specification ▪ IP Integration ▪ Verification ▪ Physical NoC Estimation ▪ RTL Generation ▪ Documentation → Physical Implementation

System IP Library

Commercially Licensed or internally developed functional IP Blocks

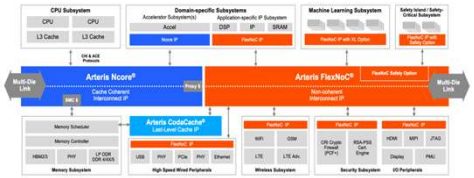

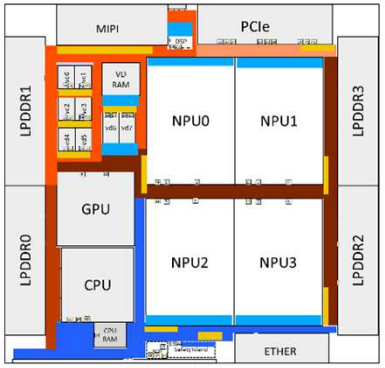

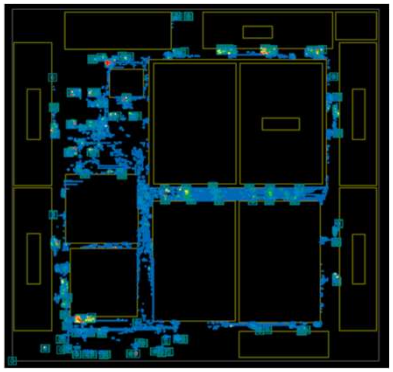

SoC Bill of Materials | Architectural View | SoC View | Physical Awareness View

*Management estimates

Need a Comprehensive System IP Foundation For Multi-Die SoCs

System IP Segments → Model, Package, Connect, Verify

SoC Integration Automation

SoC IP blocks connected, packaged & configured with SoC Integration software

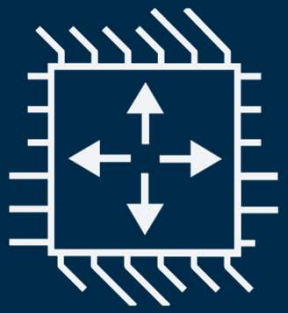

Network-on-Chip Interconnect IP

Protocol converters, switches, rate adaptors, coherent units, transport networks, directories etc.

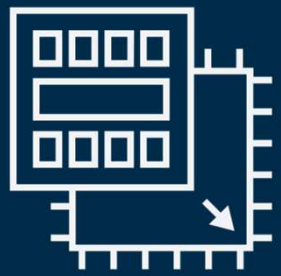

Network-on-Chip Interface IP

Additional interconnects & IP blocks connected to NoC IPs: Last Level Caches, Reorder Buffers, MMUs etc.

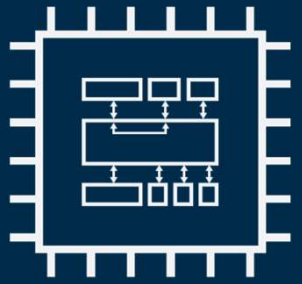

Functional safety, cyber security, physical awareness, power & domain management, scalability etc.

System IP and Network-on-Chip (NoC) SoC Interconnect IPs

Networking techniques for improved on-chip communication & data flow

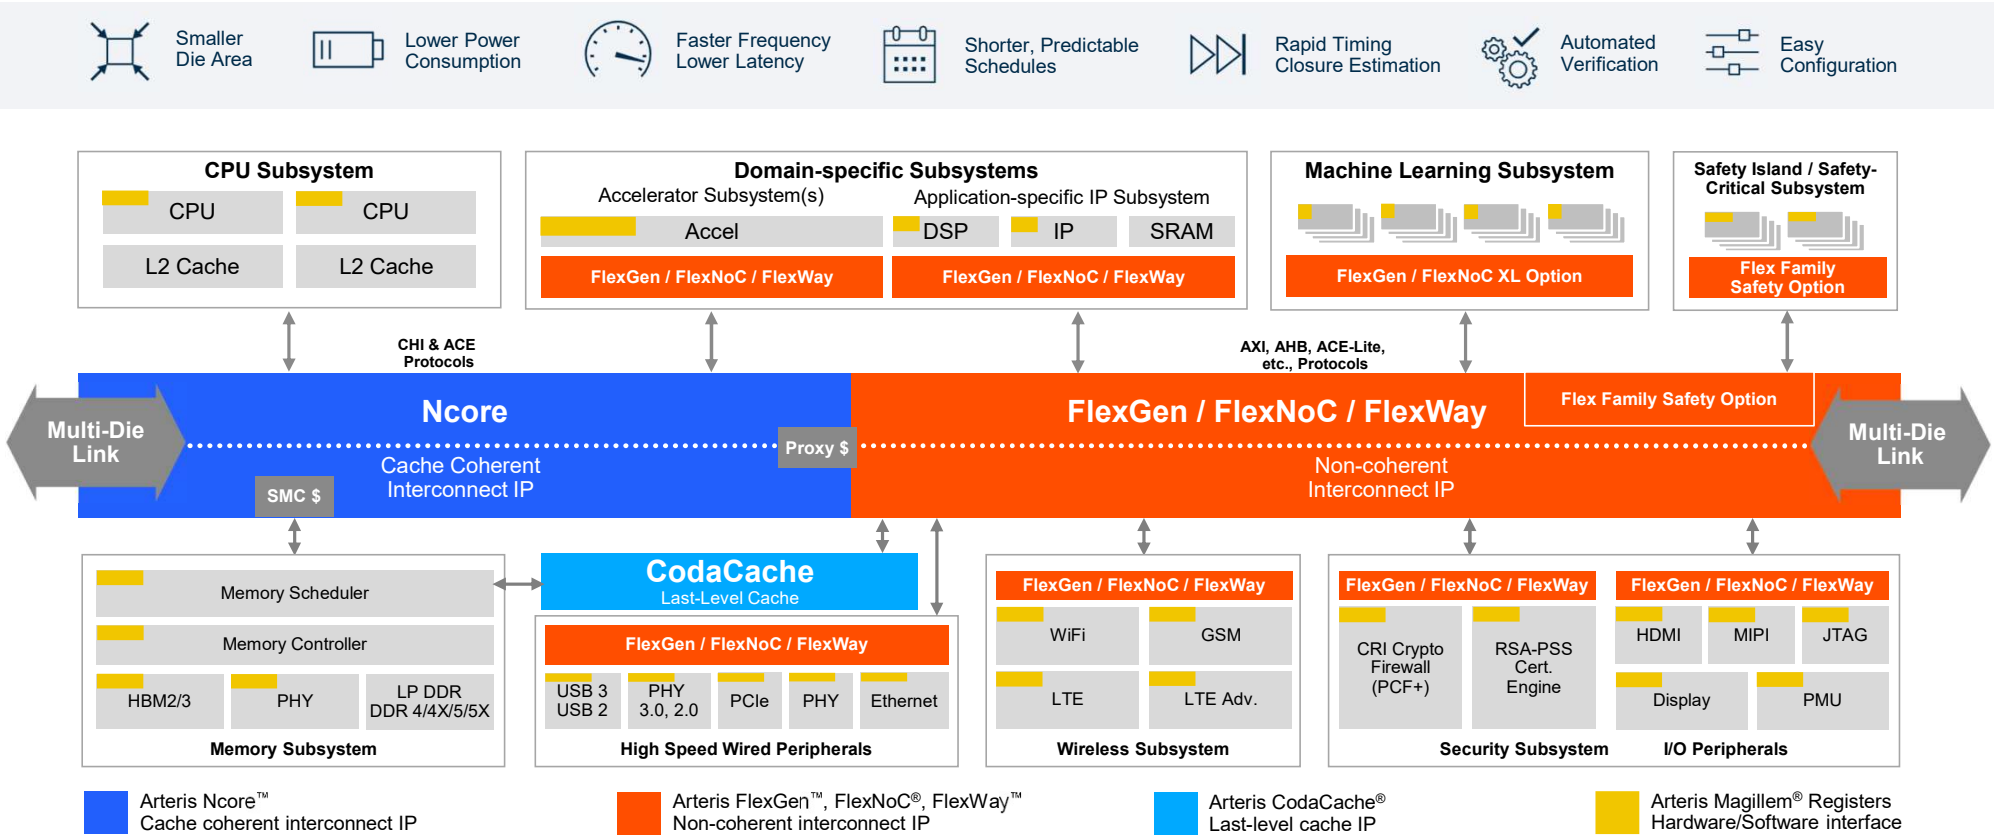

Multi-Die SoCs Increasingly Adopted For Leading Edge Applications

AI-based compute embedded in more most new designs is pushing the limits of monolithic SoCs

Requirements ▪ Specification ▪ IP Integration ▪ Verification ▪ Physical NoC Estimation ▪ RTL Generation ▪ Documentation

Physical Implementation

System IP Library

Commercially Licensed or internally developed functional IP Blocks

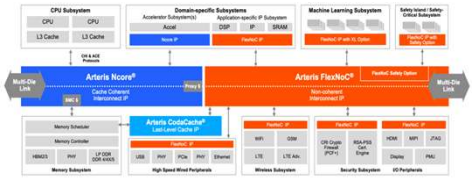

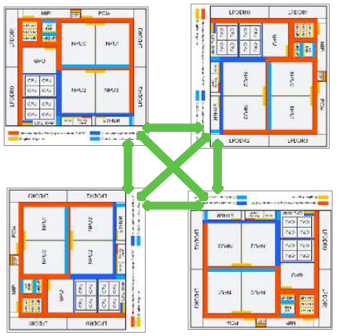

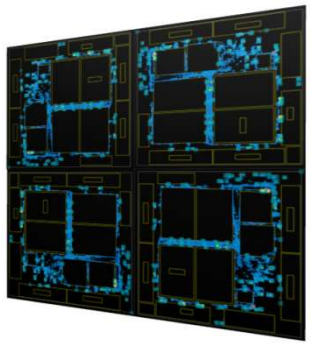

SoC Bill of Materials

Architectural View

SoC View

Physical Awareness View

*Management estimates

ARTERIS IP

Why Do Many Electronic Systems Benefit from Multi-die SoCs?

- Reticle Sizes limited to 858mm² sq → AI-base compute pushes limits
- Large Monolithic SoCs have lower yield and higher cost per chip
- Harder to re-use or modify IP quickly, vs one size fits all on same node
  - For example; XPUs on leading edge, HBM on specialized processes, AMS on trailing edge

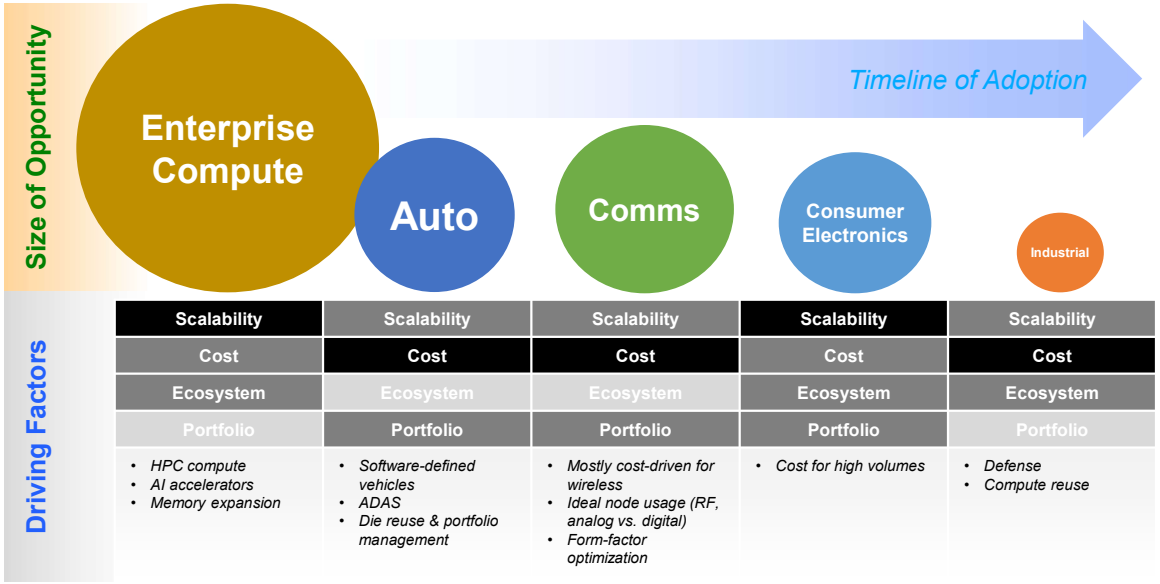

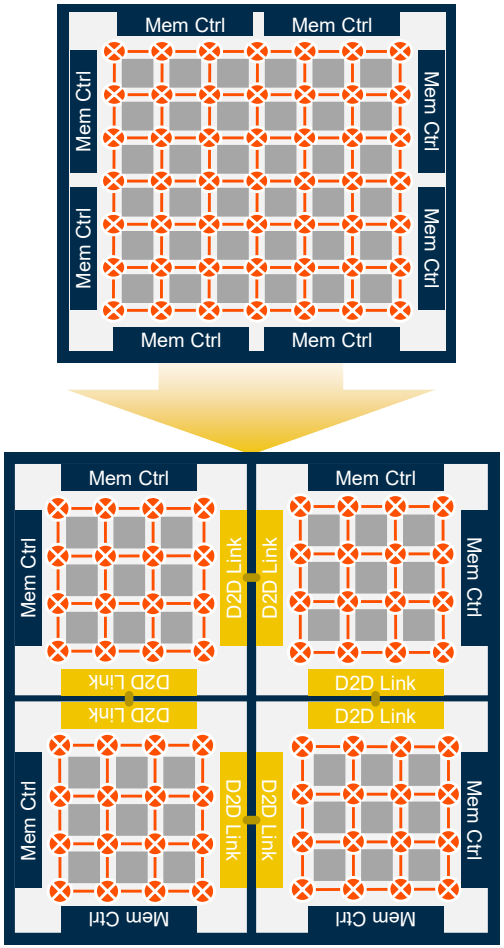

ARTERIS IP

Examples of Multi-Die Configurations

•**Non-coherent chiplets** are suitable for specialized functions where data consistency across the system is not critical or is managed by software.

•**Cache-coherent chiplets** are essential for general-purpose processing or where the shared data needs to be accessed and modified by multiple components consistently.

Non-coherent Chiplets

e.g.: AI or DSP Accelerators, GPUs, Inference NPUs, FPGAs, Networking I/O or Sensor Hubs, Vision, Cryptography, HBM I/F, etc.

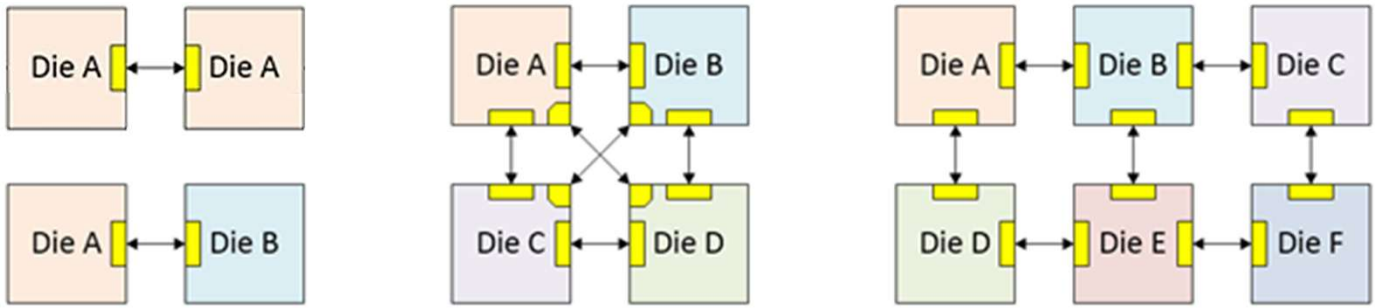

Cache-coherent Chiplets

e.g.: ADAS SoCs, CPU Compute Complex, CPU-GPU chiplets, Coherent Accelerators, Multi-Socket Servers, Virtualization, Cloud, Client Devices, etc.

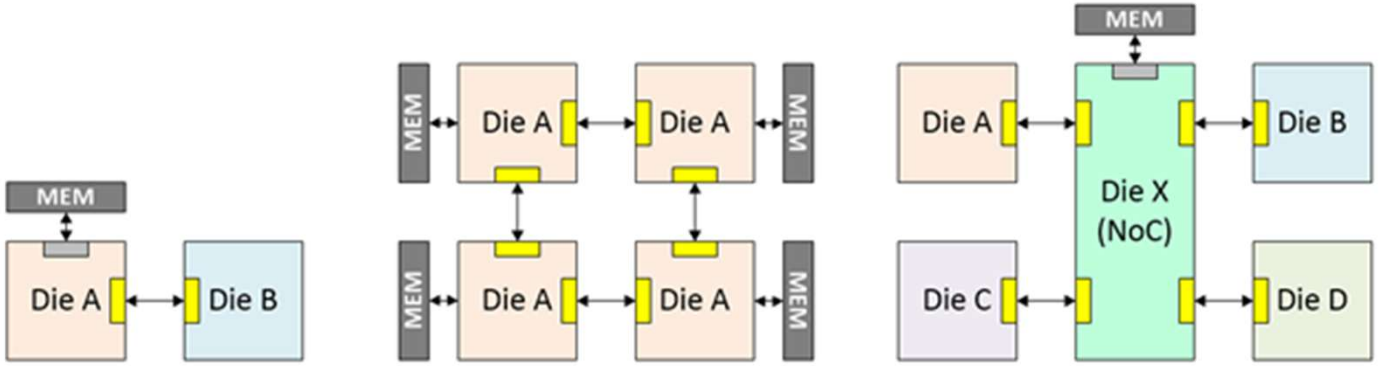

Growing complexity and data transport with bigger chiplet Scale-Out

ARTERIS IP

Progressive Complexity in Chiplet Assemblies

Homogeneous & Symmetric	Heterogeneous & Asymmetric	Heterogeneous Systems
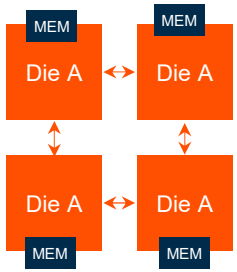	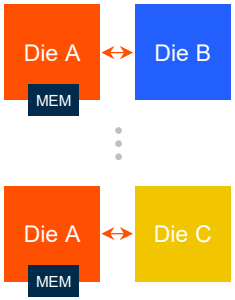	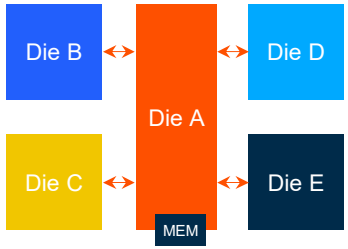
▪ Identical Chiplets ▪ Pre-planned Topologies ▪ Custom D2D Protocol	▪ Small number of Chiplets ▪ Limited mix-and-match ▪ Custom D2D Protocol	▪ IO Hub configuration ▪ Open chiplet ecosystem ▪ Interoperable D2D Protocol

ARTERIS IP

Chiplet Network Routing
Global routing is X-Y to avoid deadlock

- Chiplets are treated either as a mesh or fully-connected
- Mesh routing split into two independent tasks:
  - Local routing, on die
  - Global routing, die-to-die
- Global routing is handled by the units
  - Units are responsible for finding the next local target to reach
  - Uses a hierarchical address map decode to find its next target
- Response will not necessarily come back by the same route
  - Due to **X - Y** routing with **X** always first
  - Request: A → B → C, Response D → C → A

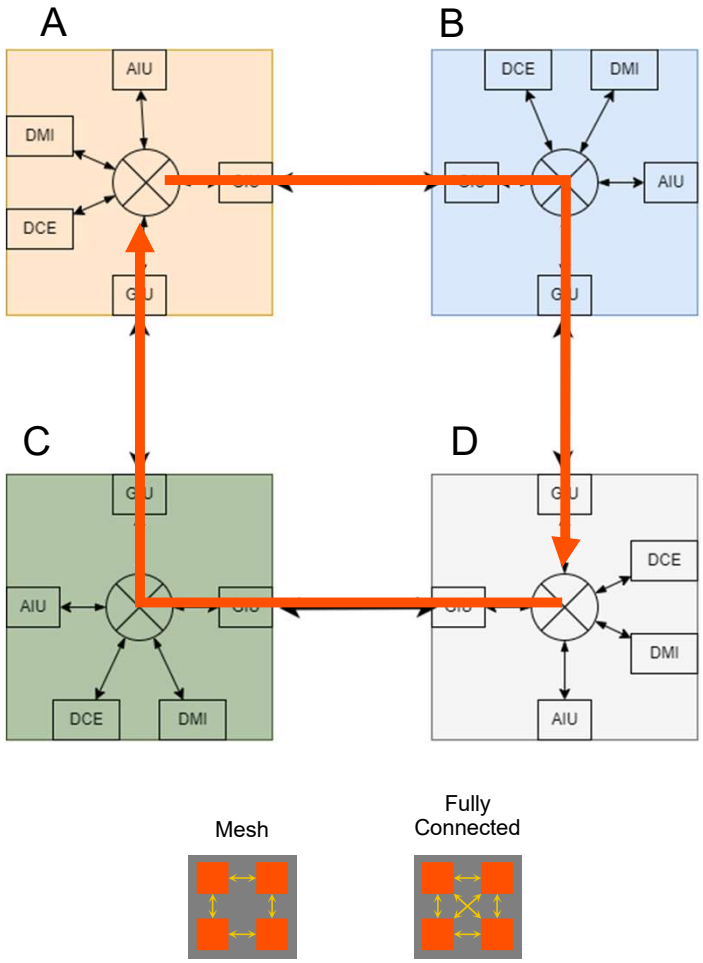

Hierarchical Address Map

- Address map is local to a chiplet
  - Defines either a local destination **or** a remote chiplet
  - Local destination will have lower latency - NUMA
- Address is re-decoded every time it crosses a chiplet boundary

Addr	Size	Dest.
A000	8kB	DMI1 Die C
8000	8kB	DMI1 Die C
6000	8kB	DMI1 Die B
4000	8kB	DMI0 Die B
2000	8kB	DMI1 Die A
0000	8kB	DMI0 Die A

System address map

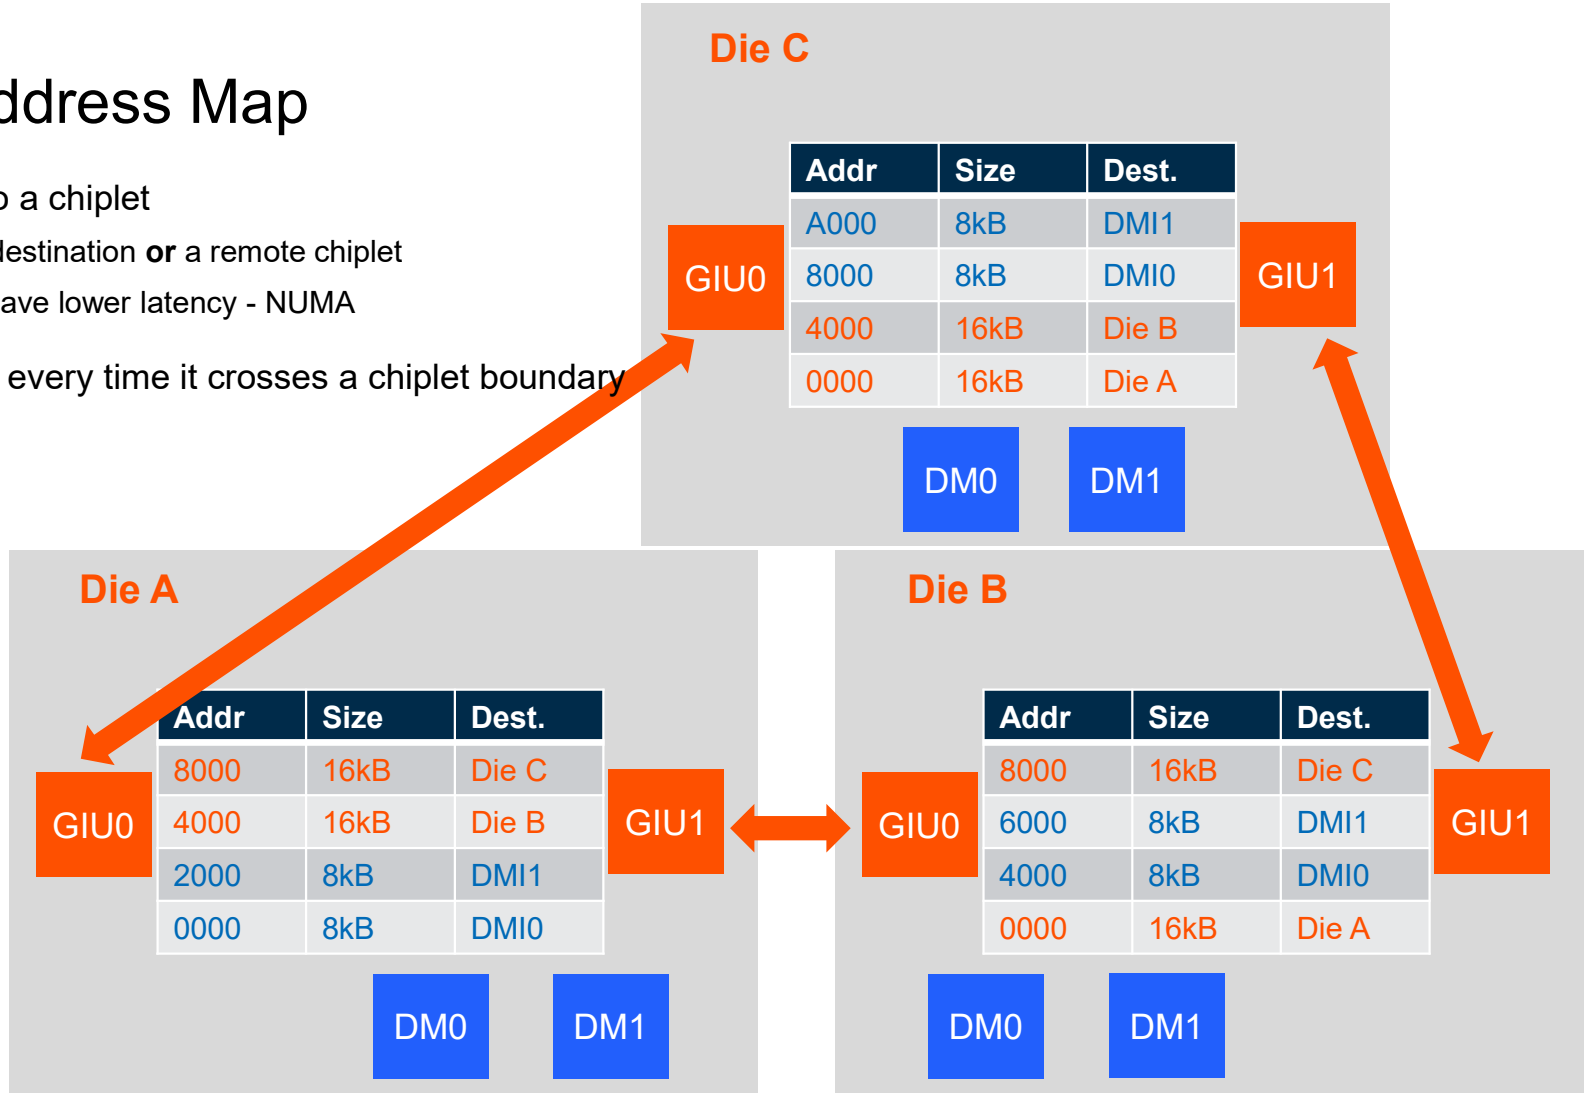

Ncore Multi-Die Fully-Coherent Die-to-Die End-to-End Solution

Making multi-die coherency look like single die coherency to software

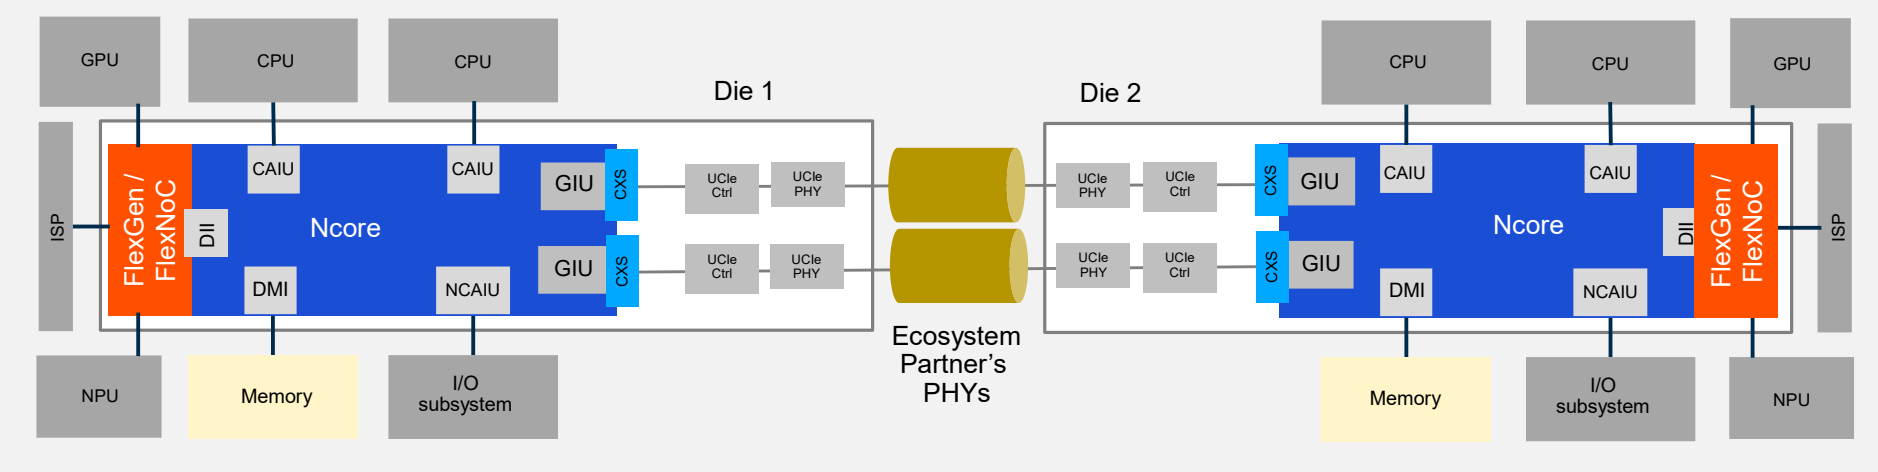

- Pre-integrated and verified die-to-die interconnect with UCIe controller and PHY
  - Support up to 4 dies and 4 links per die
  - Ncore interfaces to the UCIe controller via AMBA CXS.B interface
  - Customers will need to license 3rd-party UCIe controller and PHY for a complete solution
  - Arteris support Synopsys and Cadence initially, with additional partners expected later

Examples of Supported Coherent Topologies

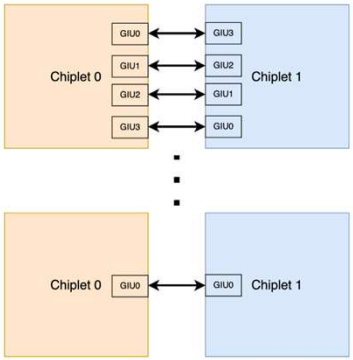

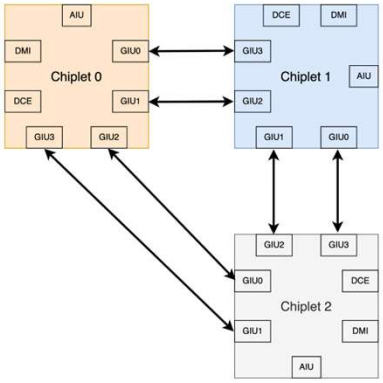

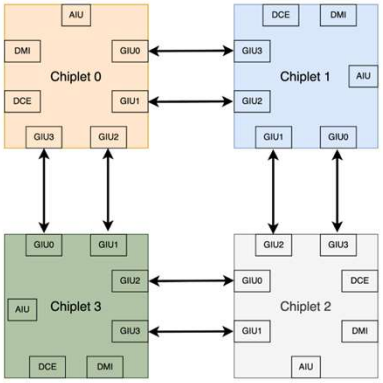

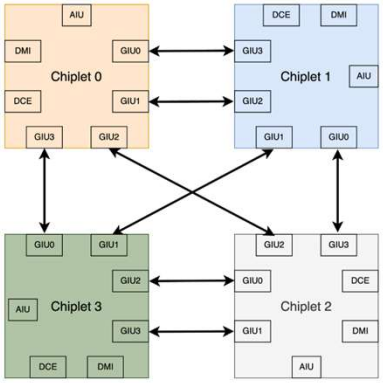

- Up to four dies/chiplets per assembly
- Up to four die-to-die links per die/chiplet

- A single coherent domain for all chiplets
- NUMA – Non-uniform memory architecture
  - Lower latency access to local memory

- One DVM MMU domain for all participating cores
  - Arm cores will likely use Distributed Virtual Memory
  - RISC-V cores typically will not use DVM

- PHY Ecosystem partner interoperability validation
- All dies are checked by Arteris tooling and verified together at design time

Partitioning Used To Disaggregate Monolithic SoC Designs Into Chiplets

Advantages

- Single source of truth IP-XACT database
- Top-down flow with disaggregation of designs into chiplets
- System Memory-Map validation with Magillem Registers
- 30% faster IP integration with rule-based IP instantiation

Magillem

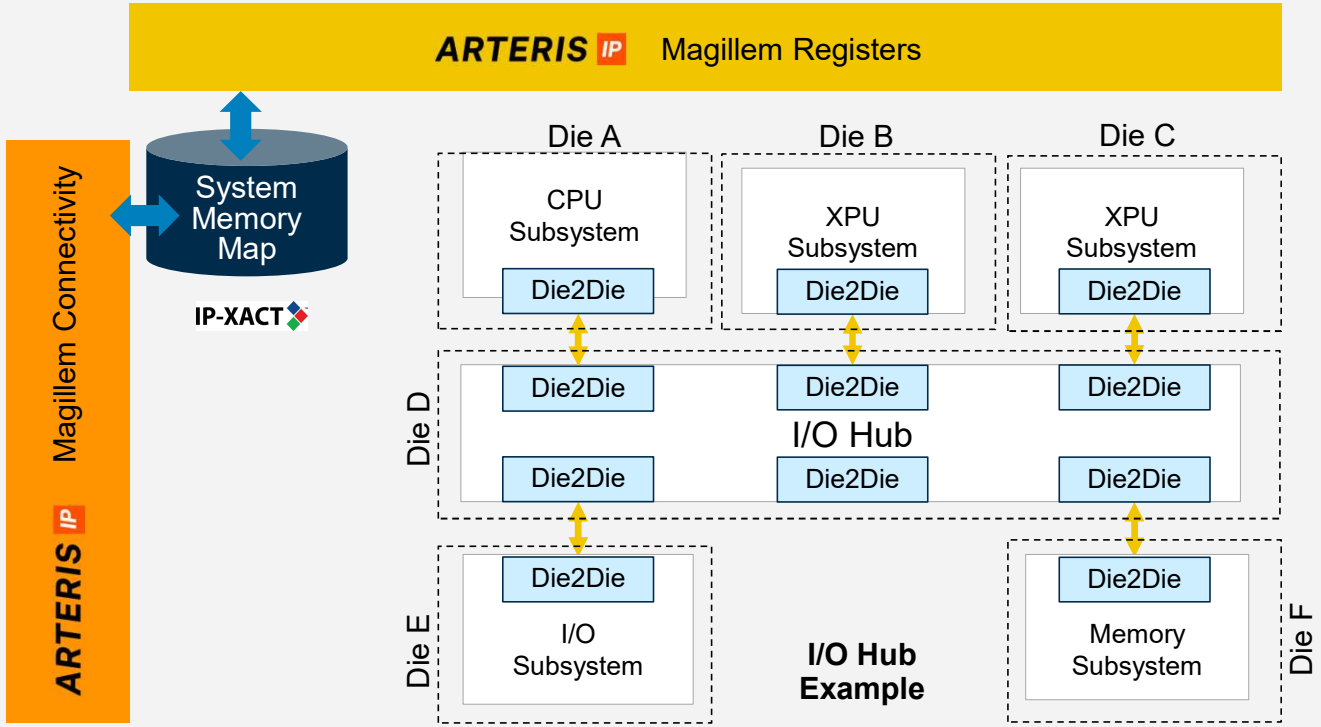

ARTERIS IP

Heterogeneous Chiplets Deal with Increased Complexity

- Chiplets used today are mostly homogeneous
- Heterogeneous chiplets are coming
- Multi-Die SoCs may communicate through central I/O Hub if More then 4 Dies or if They Are of Different Types
- There will be at multiple flavors of Interchip links
  - UCIs with AMBA AXI and/or CHI C2C, CXL, …
  - Proprietary Coherent

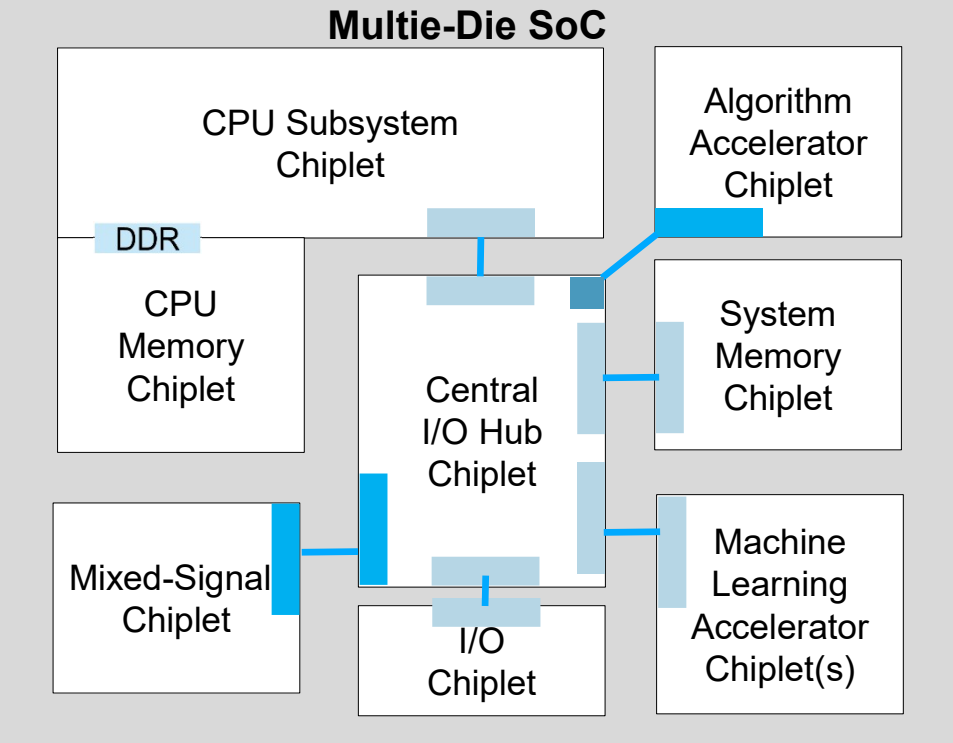

Chiplets Will Absorb More and More of Electronic Product Value

The Era of the Chiplet based SoC drives complexity and value of system IP

- 12x reticle size
- Multi-die
- Multi-vendor
- Multi-function
- Multi-process
- Multi-performance

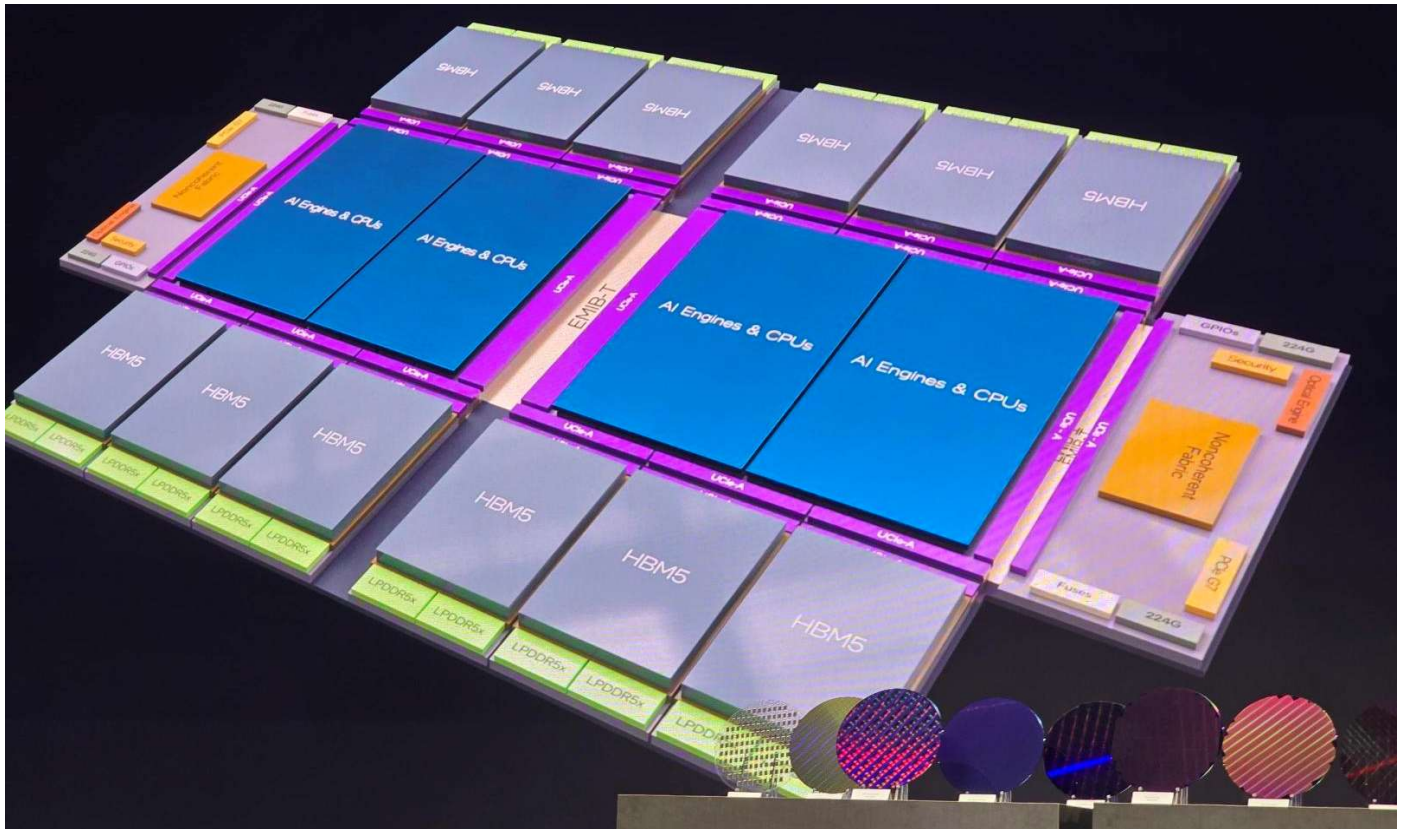

Presented by Kevin O'Buckley, SVP & GM Intel Foundry Services, Intel Foundry Day, Apr. 29, 2025

Not Just a Multi-die Solution But a Multi-die Ecosystem

System IP technology needs to be **agnostic** and connect across the **semiconductor ecosystem**, easing and derisking designs

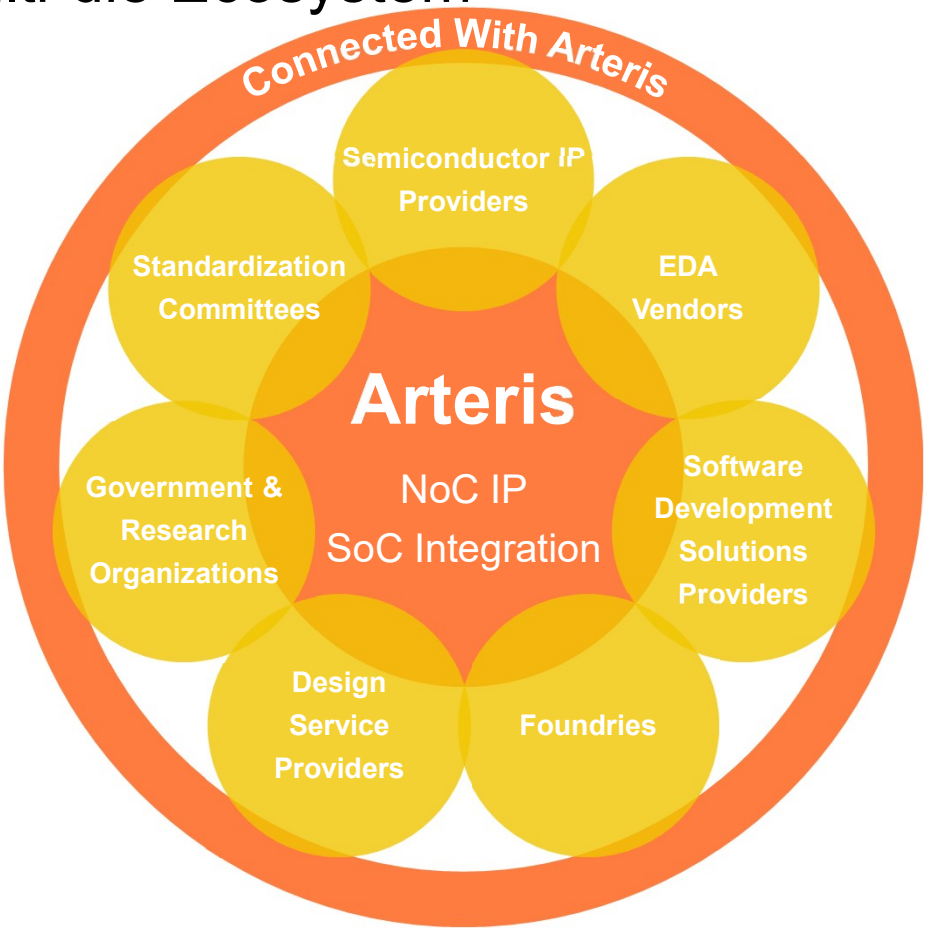

Multi-die Challenges

Many companies and universities have a role to play in the chiplet economy

- Large verification spaces → ensuring quality, robustness and reliability
  - Functional safety, cyber security, fault isolation
- Standards adherence and compatibility
  - Is everyone's UCIe and IP-XACT the same? → Need conversion flexibility, vendor cooperation, VIPs
- Multi-die Floorplanning
  - Physical effects impact Multi-die SoC performance and even functionality
- Heterogeneous cache coherency is challenging
  - Arm CMN working with Arteris Ncore? → Non-coherent maybe safer in the short term
- Cost → Today while net positive on HPC designs, Multi-Die costs are still prohibitive for most
- Heterogenous will take 2-4 years to work out all of the issues → Ecosystem cooperation, pioneering customers, multiple technology delivery, but solid benefits will be there

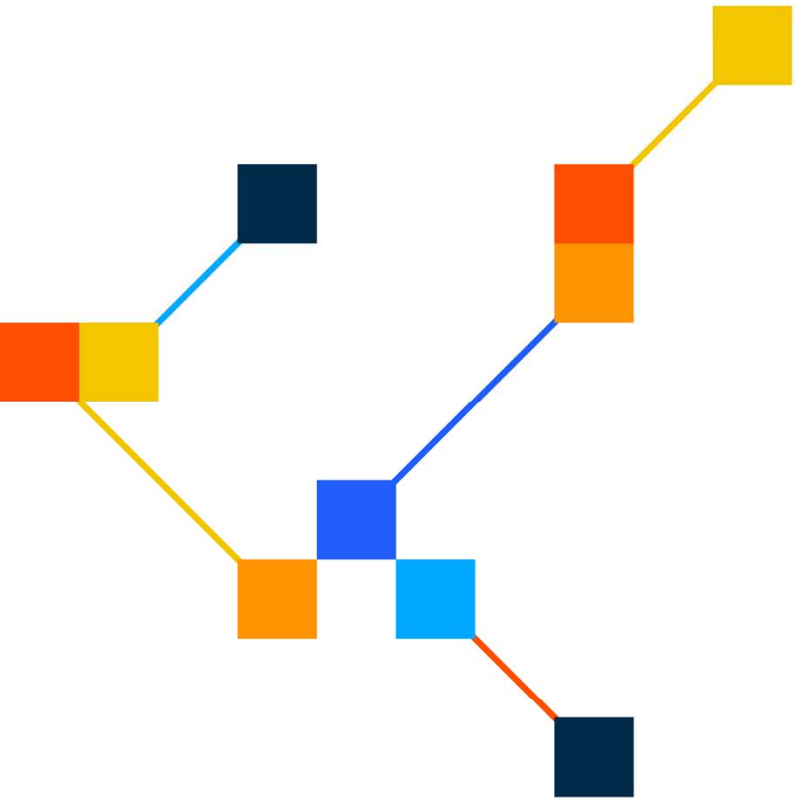

ARTERIS IP

Thank you