Single Chip and Multi Chip Integration

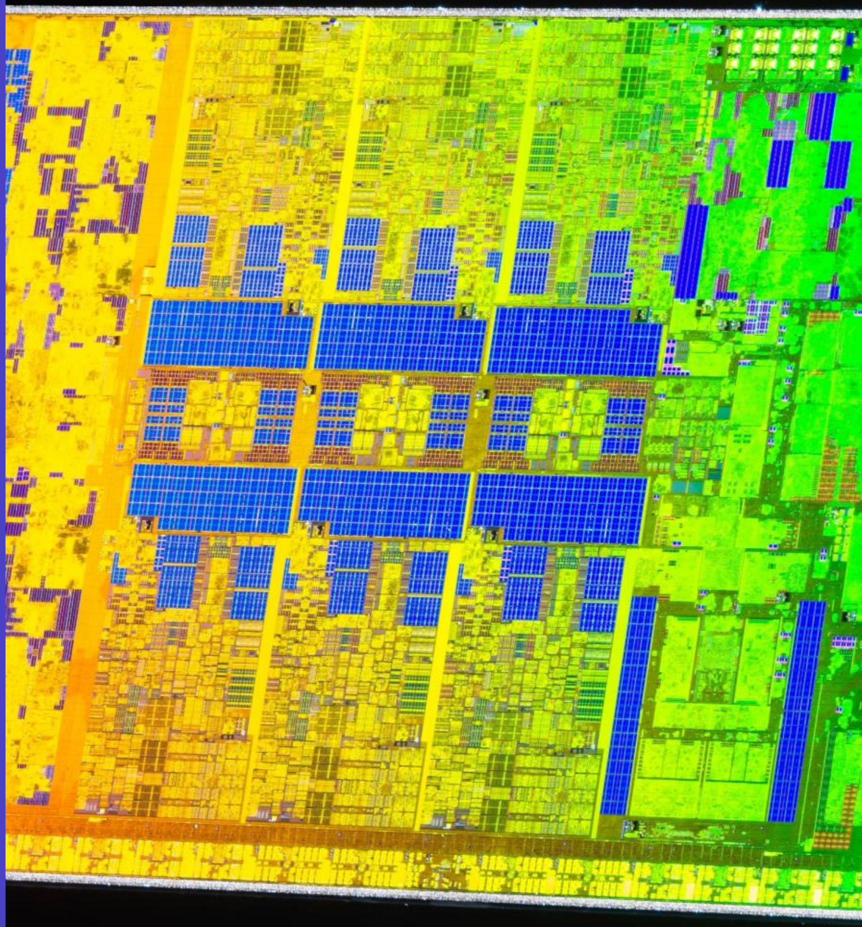

IEEE EPS @ RPI

David King

General Updates

IEEE EPS has been ACTIVE!

- Mercer XLab Renovation Grand Opening
    - Met with Industry and Academia members
    - Establish critical connections with faculty to gain support
    - Promote student involvement and other organizations.

IEEE EPS has been ACTIVE! (cont.)

- Collaboration with partner orgs
  - Industry Night with Global Foundries
    - *AiChE*
  - Dennis Group Info Session
    - *SWE*
- Albany Nanotech Complex Tour
  - Visited Nation's Largest Semiconductor R&D Facility
  - Meeting with IEEE EPS professionals
- <u>2nd tour date: November 23rd</u>

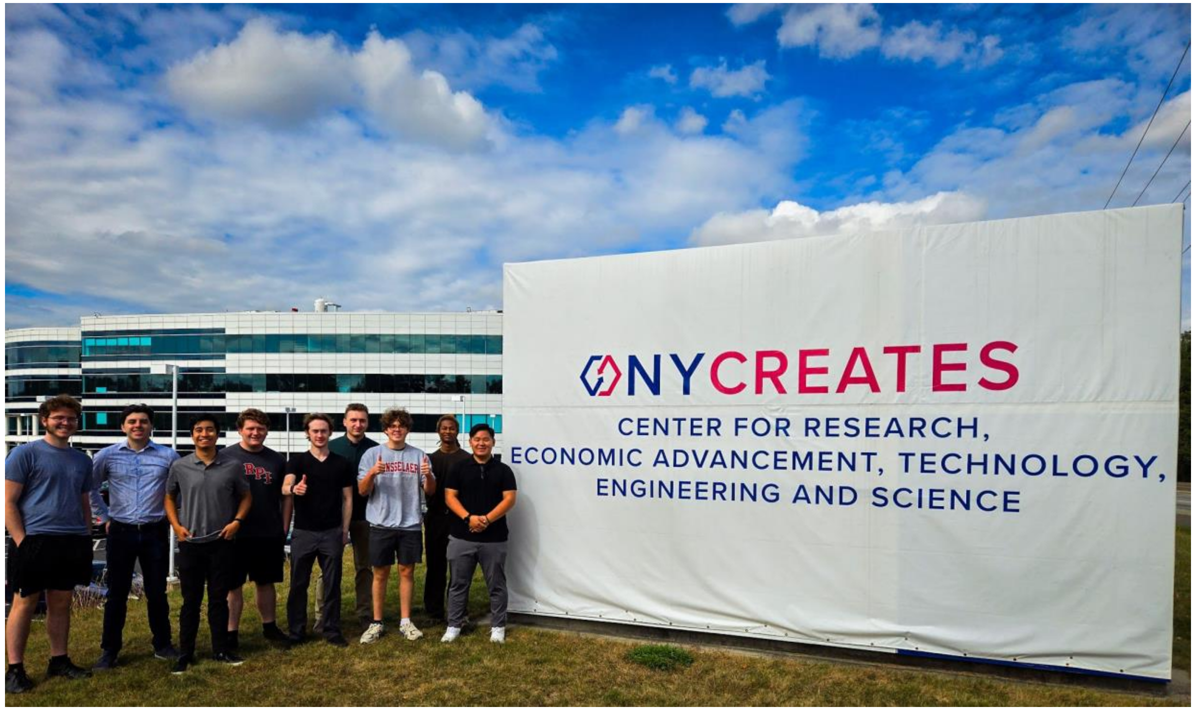

ITHERM Challenge!

Teams of <5

Open Workshop (Time TBD)

Fast paced development: join to learn OR join to compete!

2024-25 ASME/K-16 and IEEE/EPS Student Design Challenge:

Expand the Possibilities of Cold Plate Design using Additive Manufacturing

LIQUID COOLED COLD PLATES

3D PRINTED **CFD SIMULATION**

Registration: https://forms.gle/fcPV343qBXXmy4dj6
Additional Competition Information: https://bit.ly/CPD2025
Questions: ASMEHeatSinkChallenge@gmail.com

SCHEDULE

Design Guidelines Release
September 23, 2024

White Paper & Design Model Due
December 20, 2024

Semifinalist Design Revisions
January 10 – 24, 2025

Printing and Testing
January 27 – April 30, 2025

Finalist Announcement
April 30, 2025

Finalist Presentations at ITherm 2025
May 30, 2025

ITherm
DALLAS, TX
2025

Competition Leadership Committee
Prabhakar Subrahmanyam (Intel) Sameer Rao (Utah)
Chirag Kharangate (Case Western Reserve U.)
Han Hu (U. Arkansas) Tiwei Wei (Purdue University)

Competition Sponsor:
FABRIC8LABS

What are we learning today?

- The End of An Era
- A Solution: The Chiplet
- Chip Communication Basics
- The Chiplet Interconnect Ecosystem
- Challenges

Integrated Circuit Innovations

Moore's Law	Continuous growth in transistor count
Dennard Scaling	Power per transistor shrinks as speed and density increase Energy expended per computation is reducing
Clock Rate Evolution	3 MHz to 4 GHz (improved through both technology and architecture)

1.4x Annual Performance Improvement for 40+ Years ≈ 10,000x

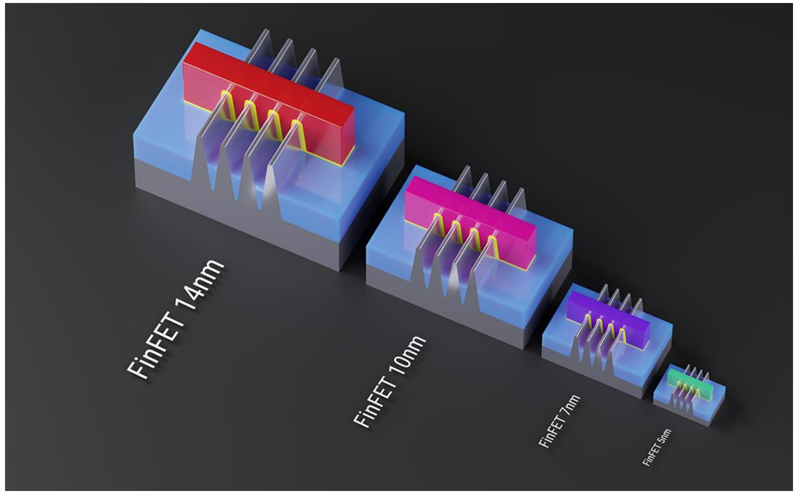

Uni-Processor Performance (Single Core)

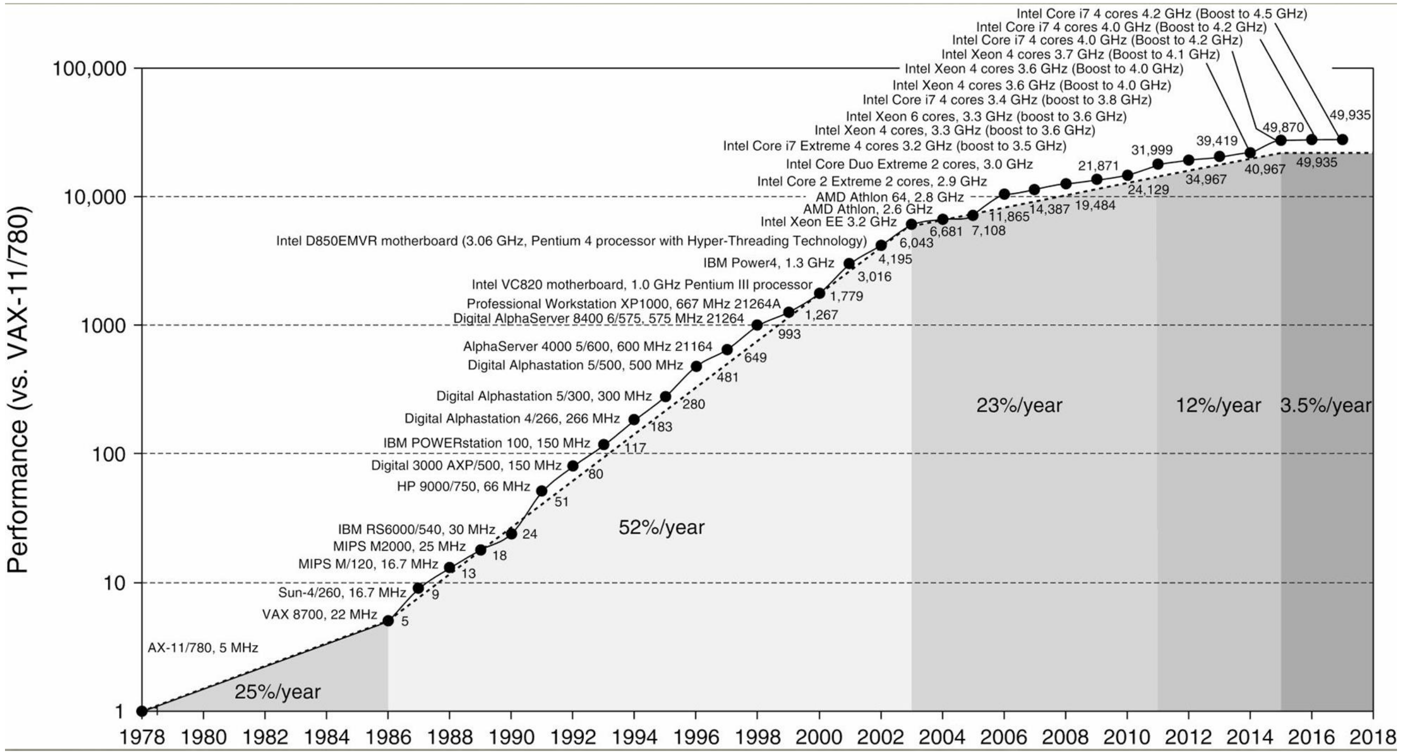

Plateauing of DRAM Density, and Computing Performance. Source: J Hennessy, ERI Conf July 2018

Architectural Innovations

Width

- 8-bit -> 16-bit -> 64-bit (~4x increase)

Instruction-Level Parallelism (ILP)

- 4 to 10 cycles per instruction
- Now, 4+ instructions per cycle (~10-20x improvement)

Multicore

- One processor to 32 cores (~32x increase in throughput)
- 6x faster throughput!

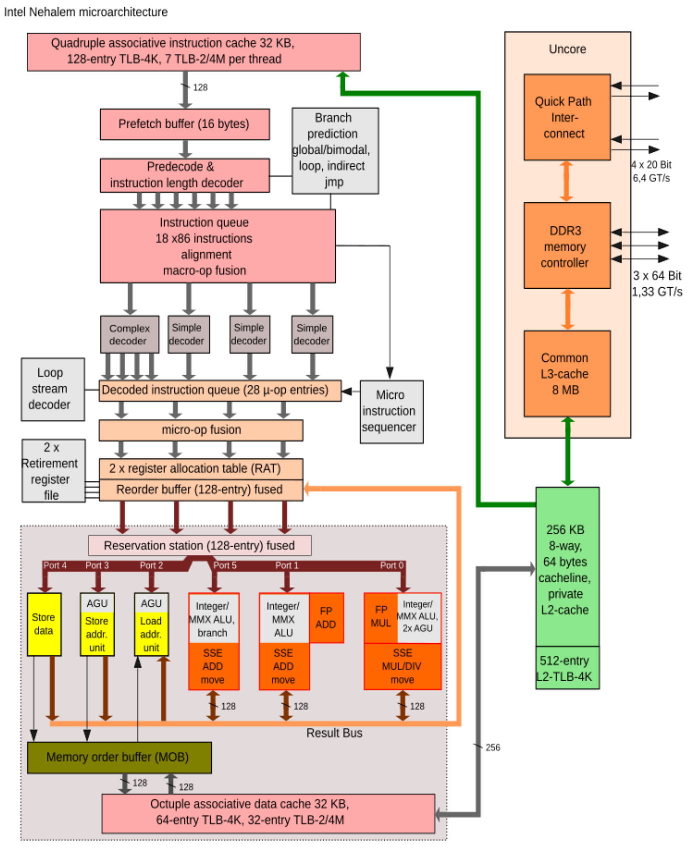

GT/s: gigatransfers per second

Three Changes Converge

IC Technology

End of Dennard scaling: power becomes the key constraint

Slowdown in Moore’s Law: transistors cost (even unused)

Architectural

Limitation and inefficiencies in exploiting instruction level parallelism end the uniprocessor era.

Application focus shifts

Amdahl’s Law and its implications end the “easy” multicore era

From desktop to individual, mobile devices and ultrascale cloud computing, new constraints

The Solution?

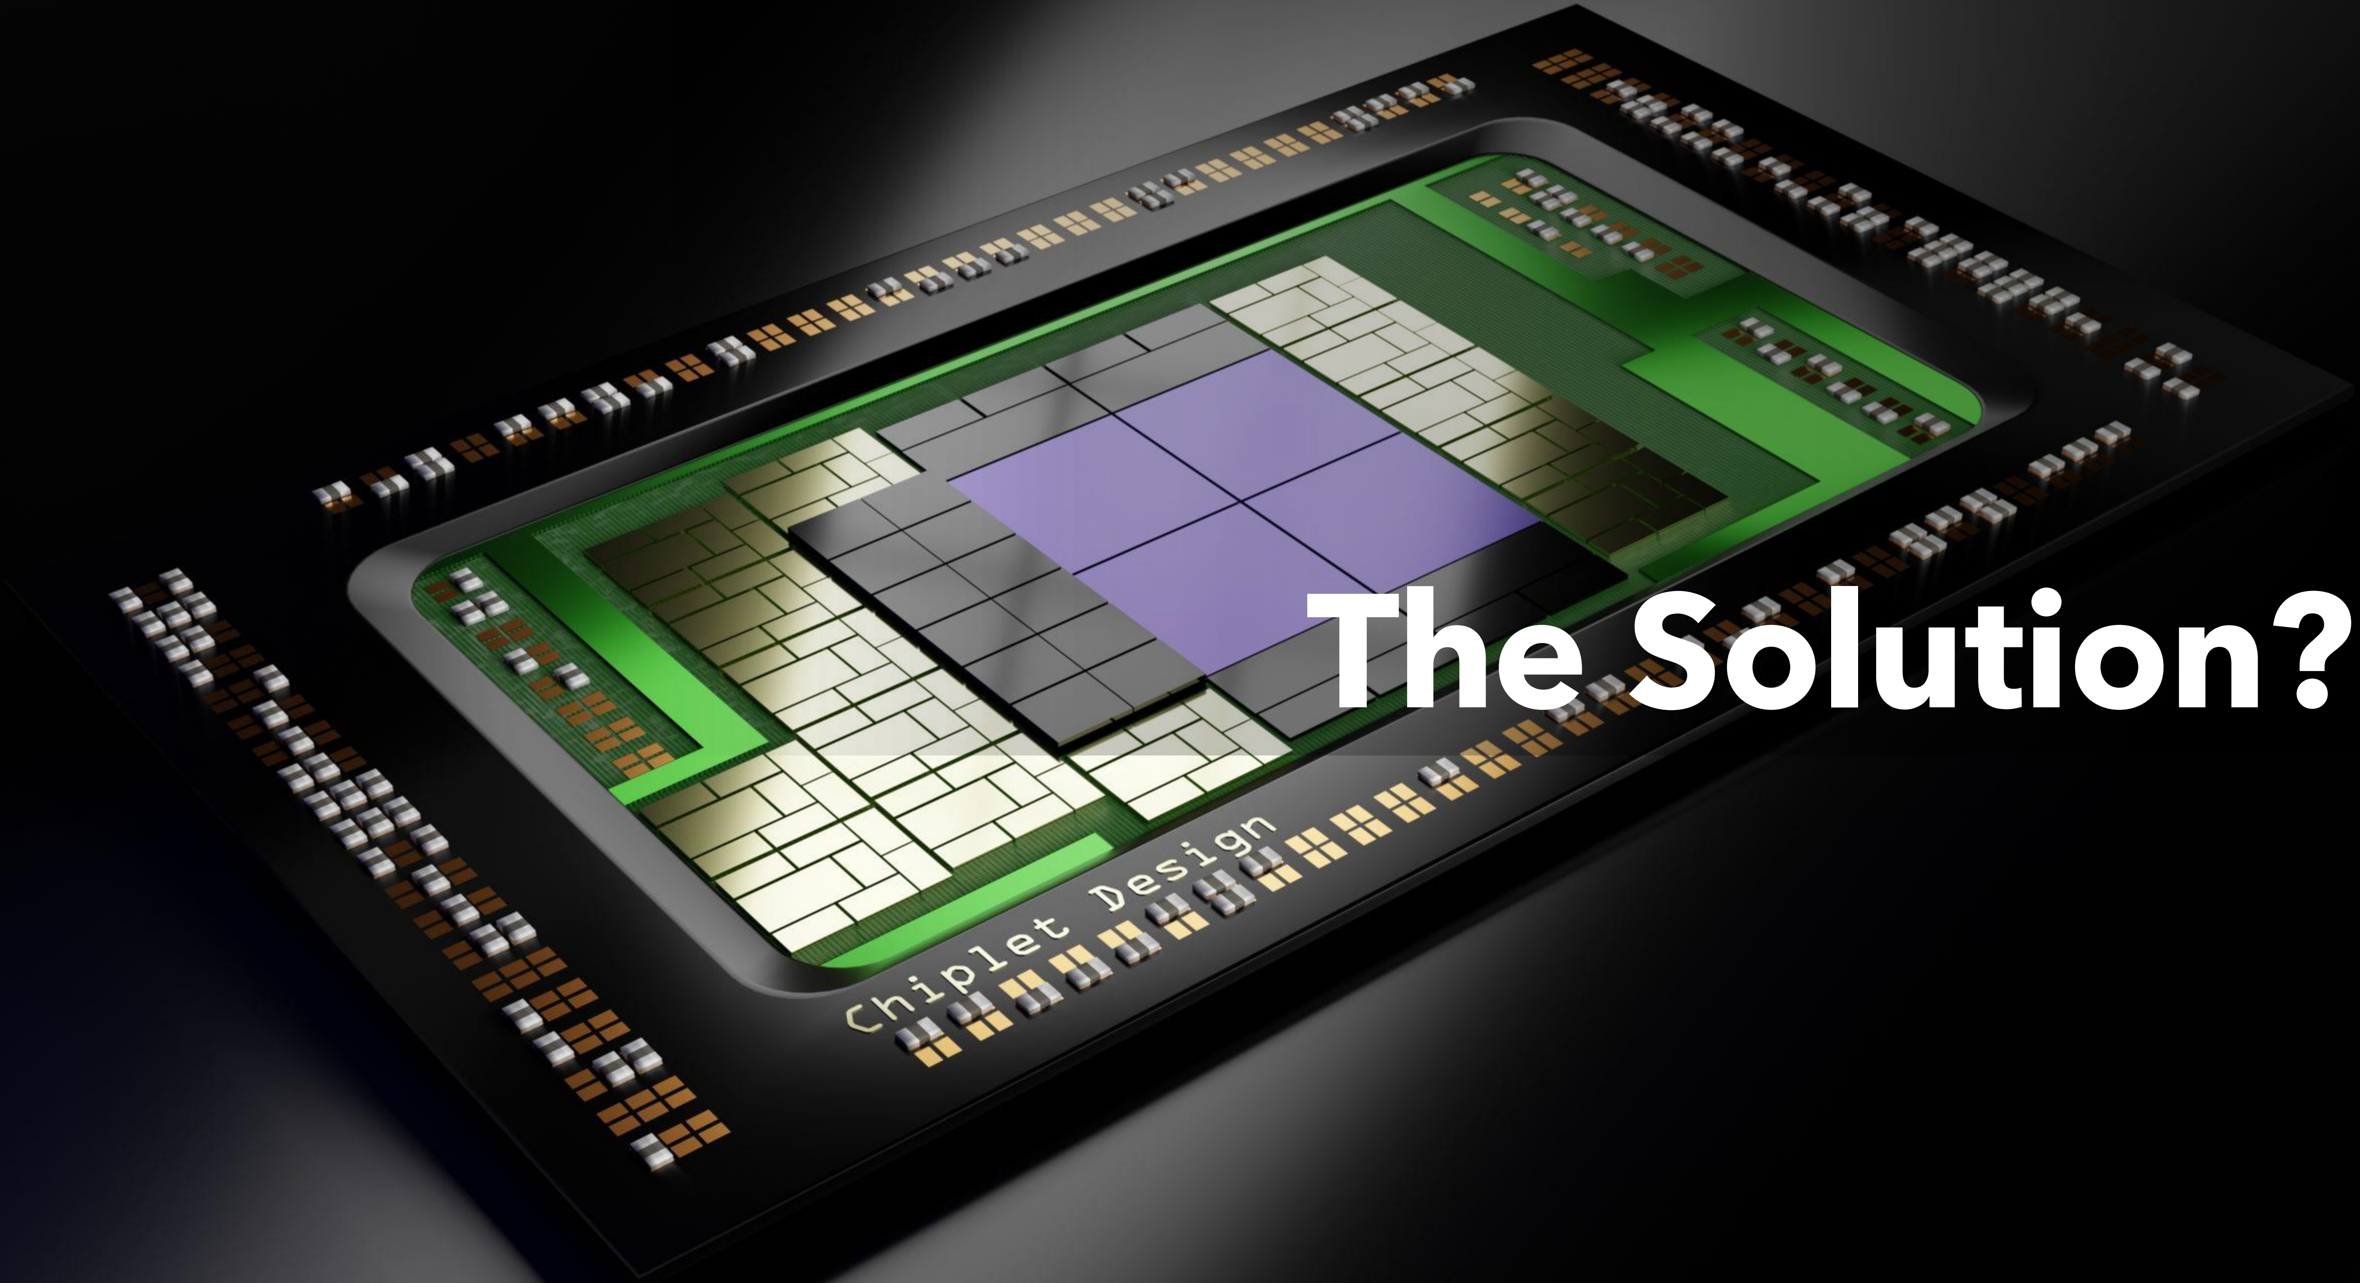

Following Moore's advice even today

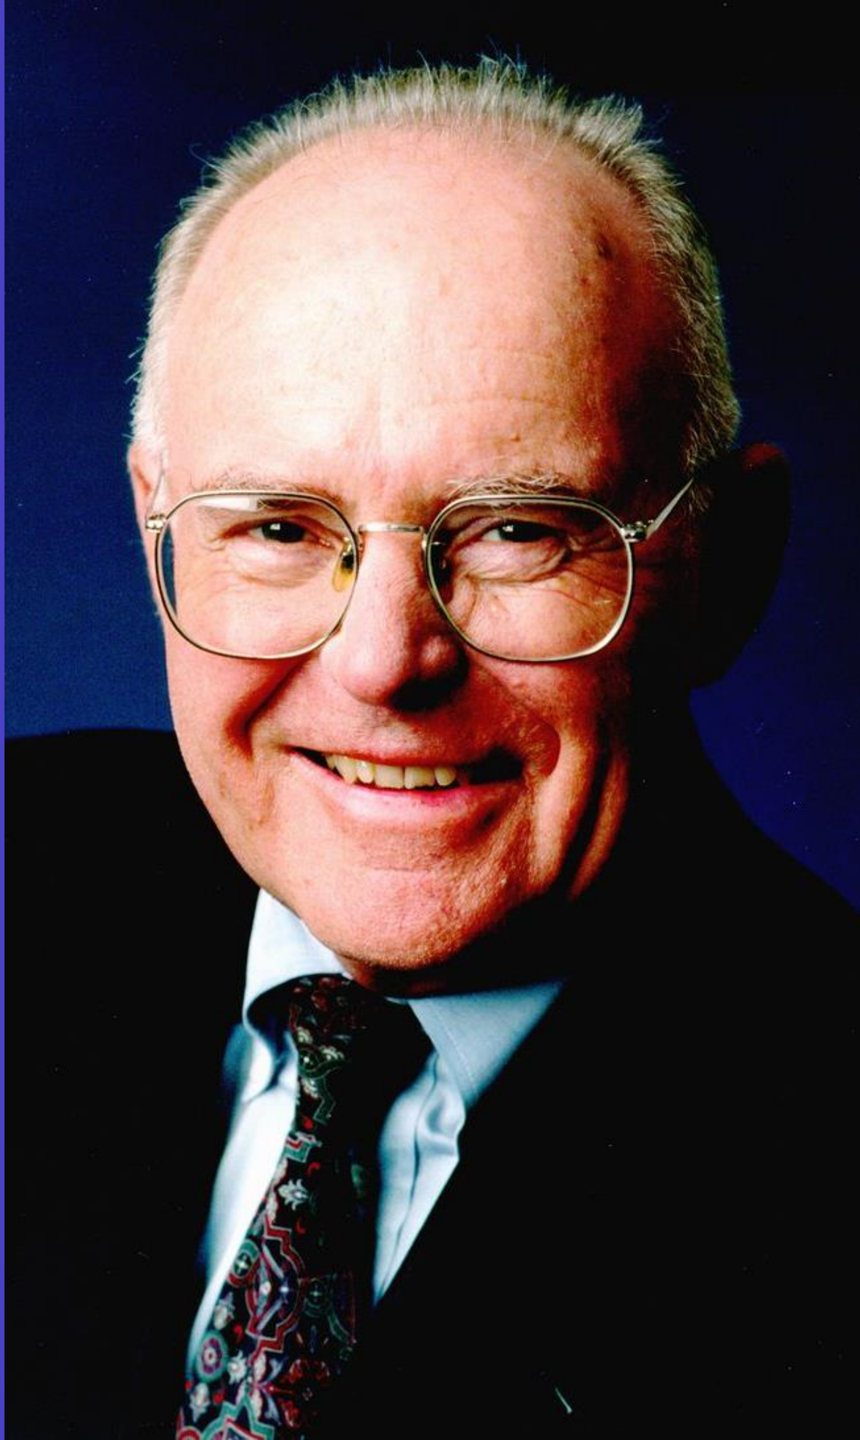

'It may prove to be more economical to build large systems out of smaller functions which are separately packaged and interconnected.' – Gordon Moore

General thought process for chiplets

Primary:

- Moore's law slowing

Secondary:

- Economics
- Flexibility
- Scalability

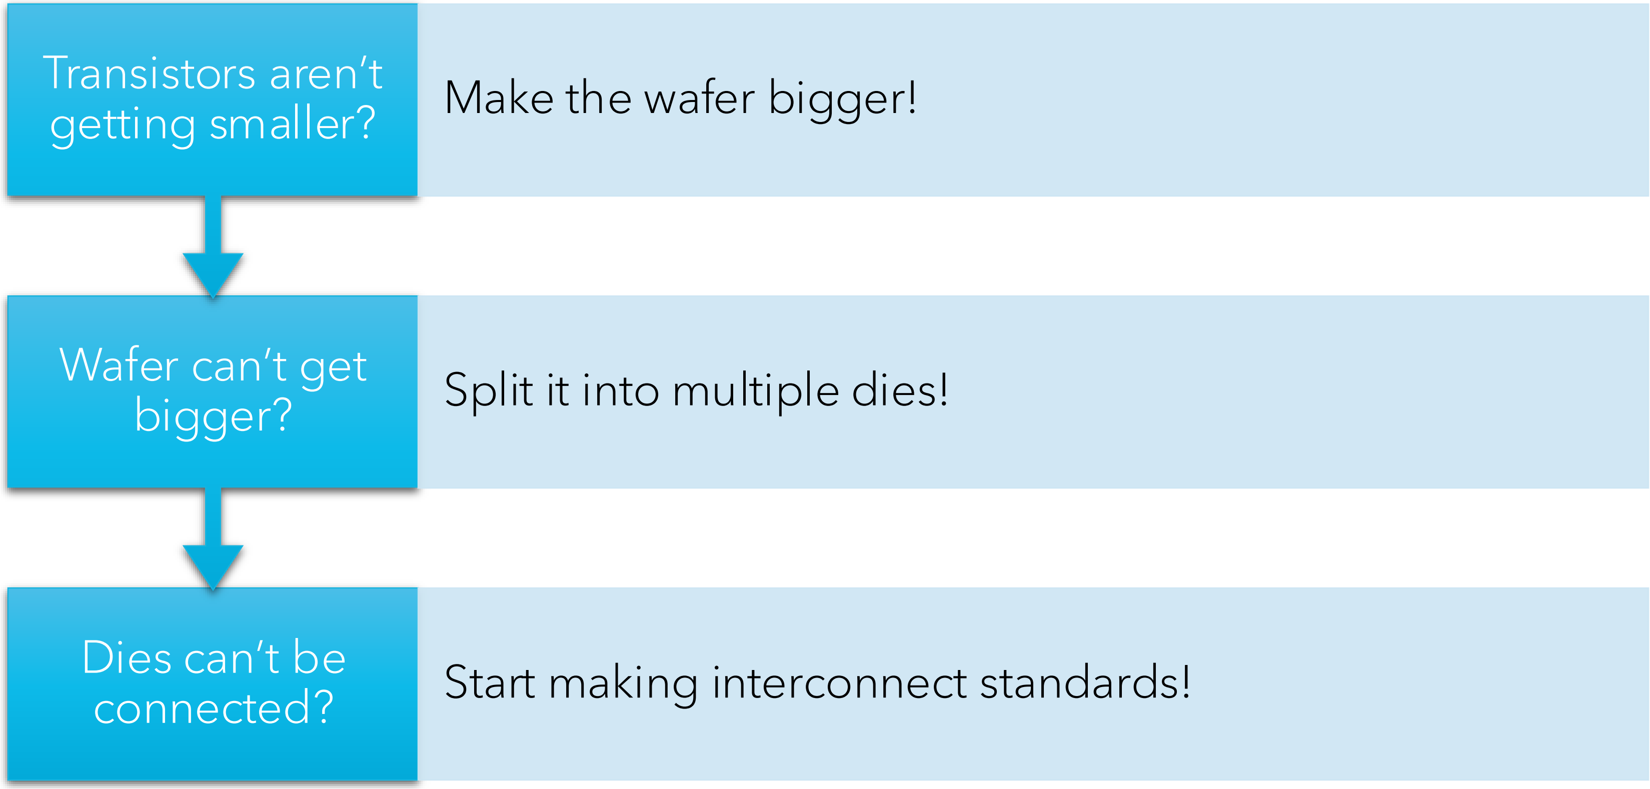

Transistors aren't getting smaller?	Make the wafer bigger!
Wafer can't get bigger?	Split it into multiple dies!
Dies can't be connected?	Start making interconnect standards!

It’s all about the interconnect

Heterogeneous integration and 3D integration depend entirely on good interconnect:

1. Standards
2. Manufacturing

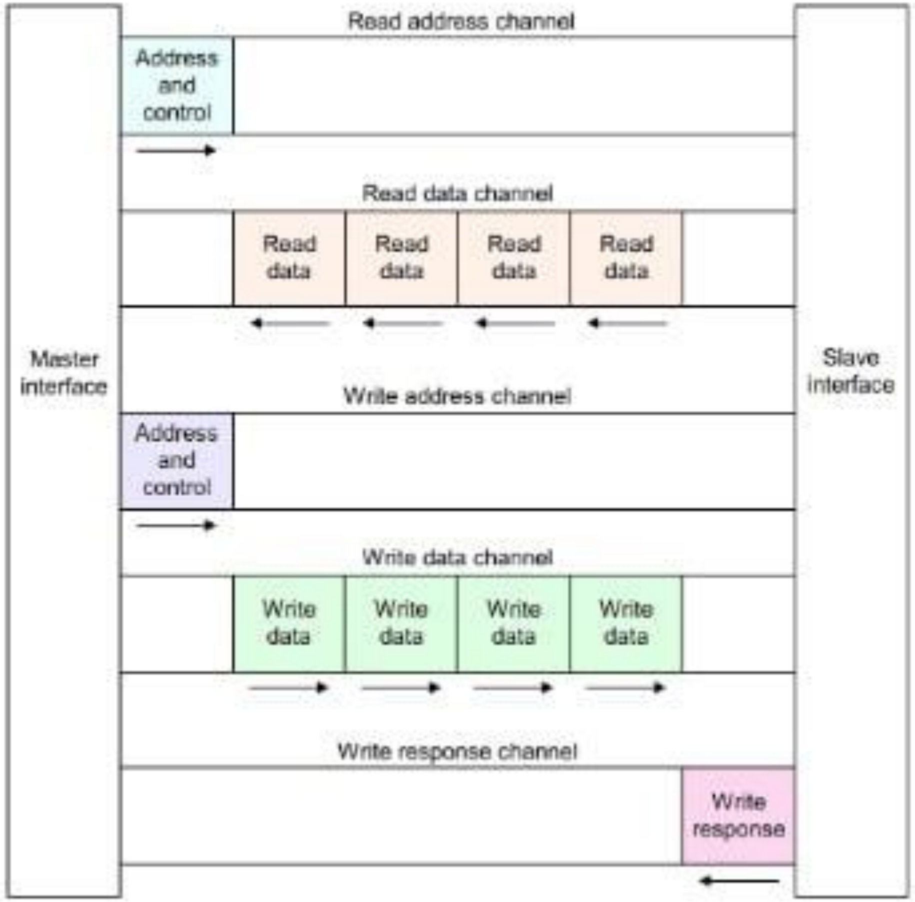

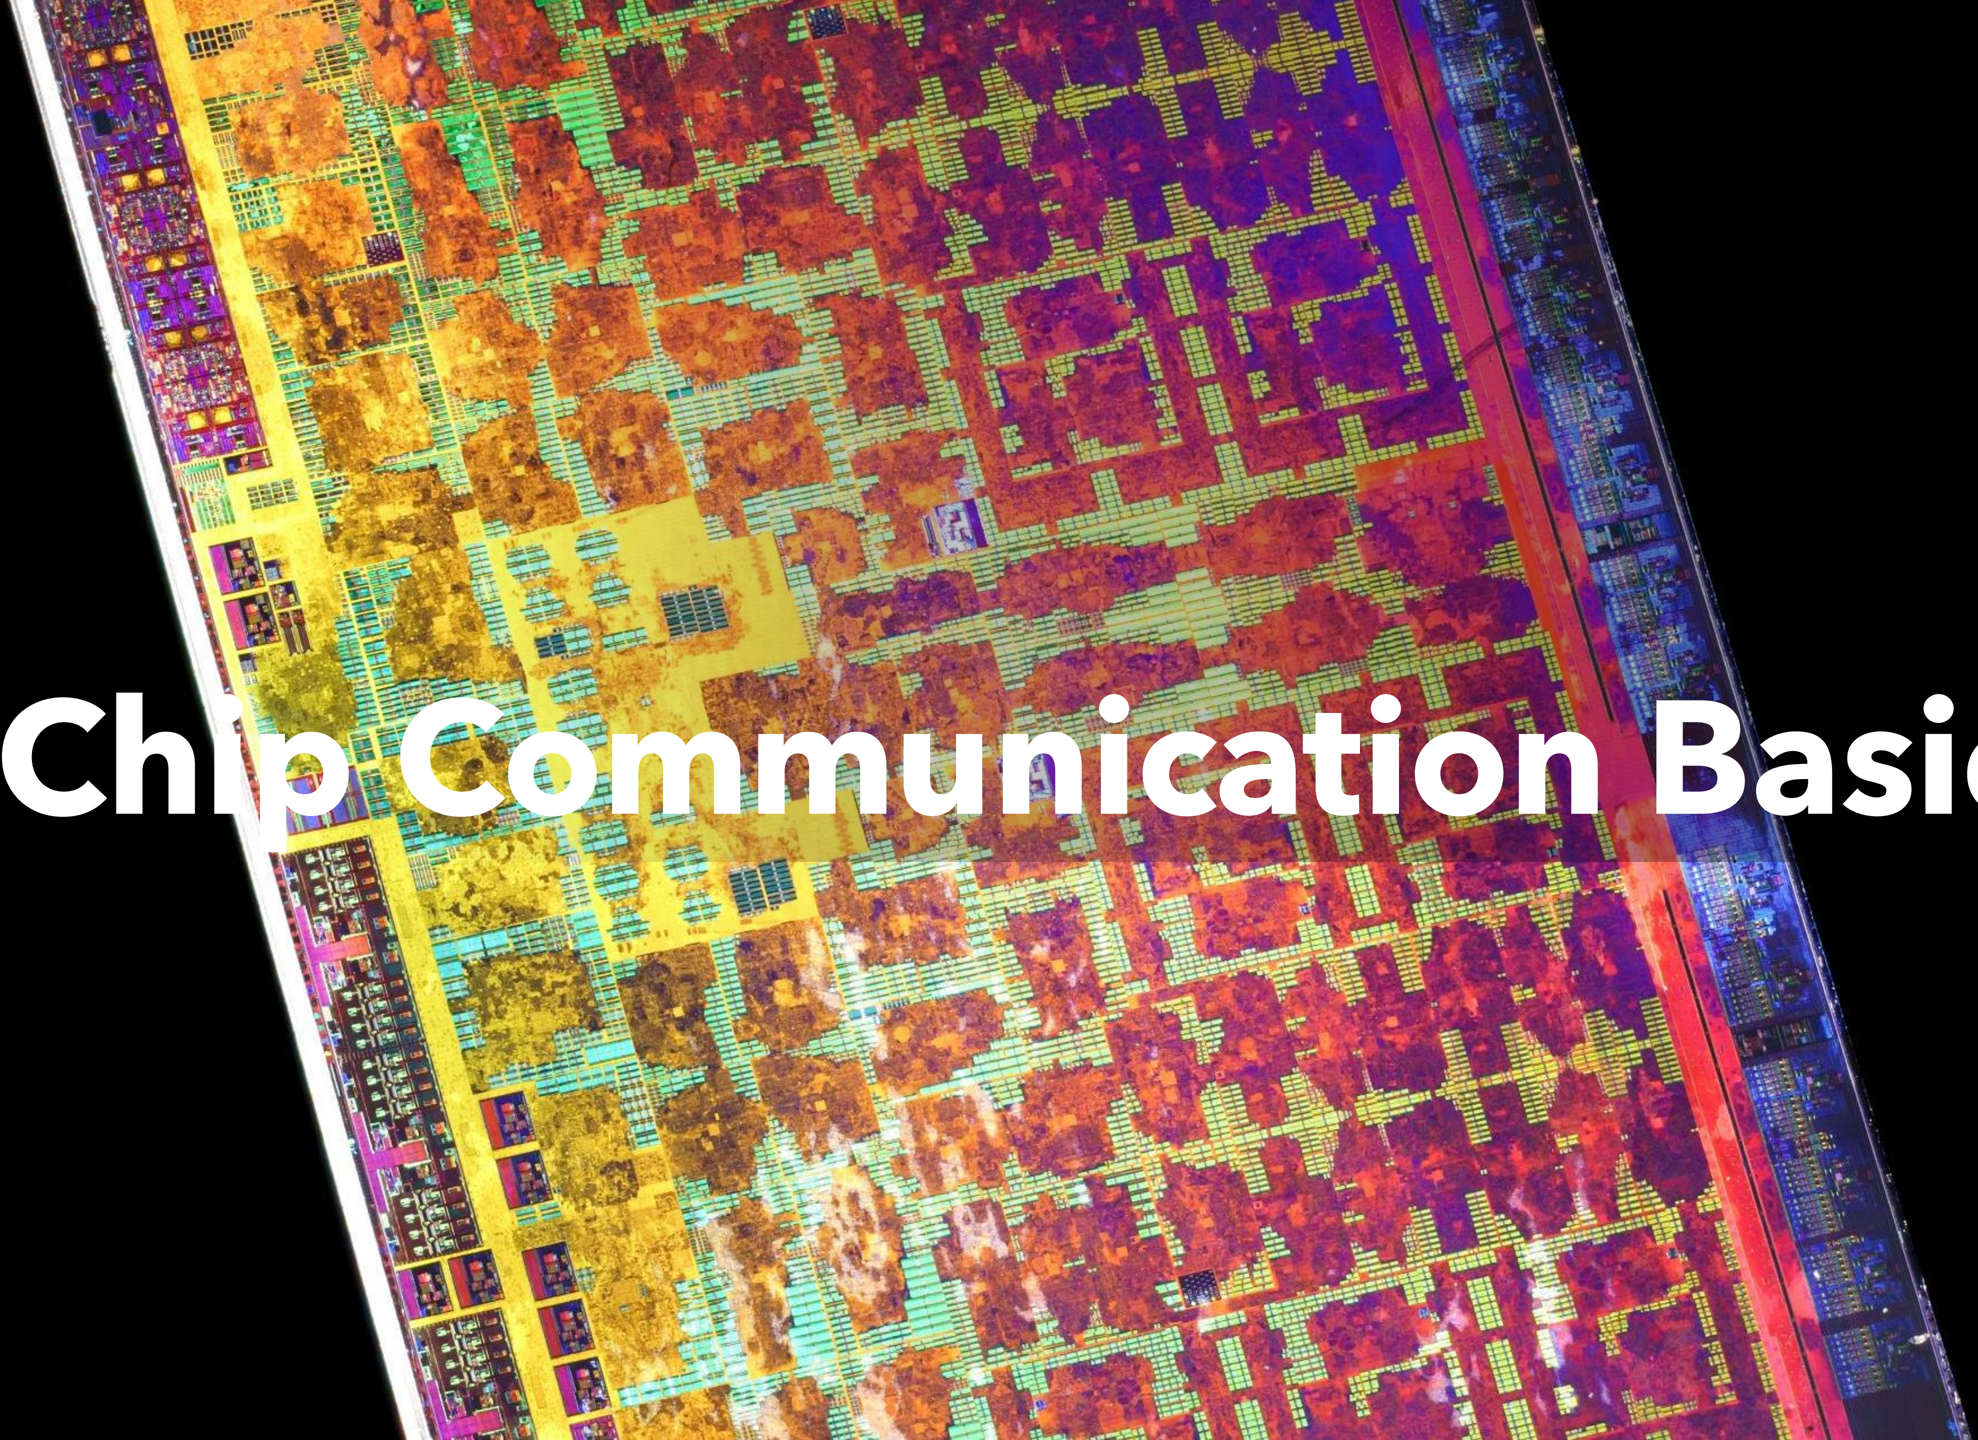

Chip Communication Basics

Building Abstraction

Application	Application oriented/specific
Transport / Transaction Layer	Ensures correct data transmission from end to end.
Data Link Layer	Ensures correct data transmission from device to device.
Physical (PHY)	Represents physical implementation of interconnect

- Productivity and reusability come with abstraction.
- Abstraction will come from:
  - Structure
  - Protocol
  - Network

Physical Layer

- The physical circuit and medium that transfers information between devices
- Functionalities:
  - EMI management
  - Signal Losses
  - Error Correction

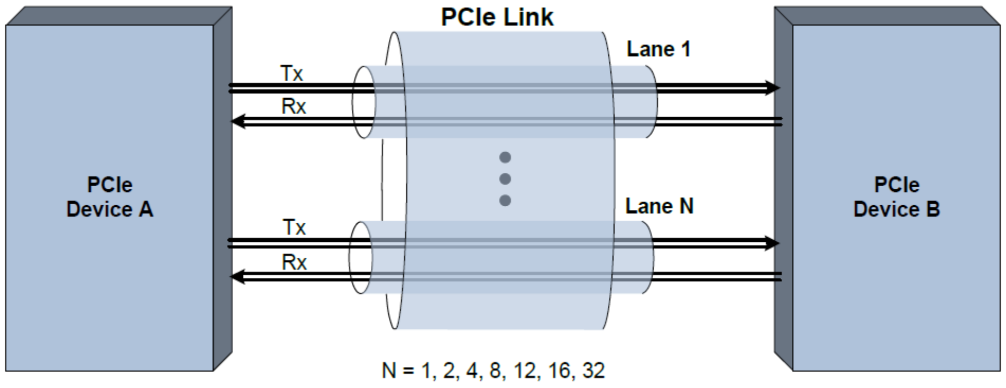

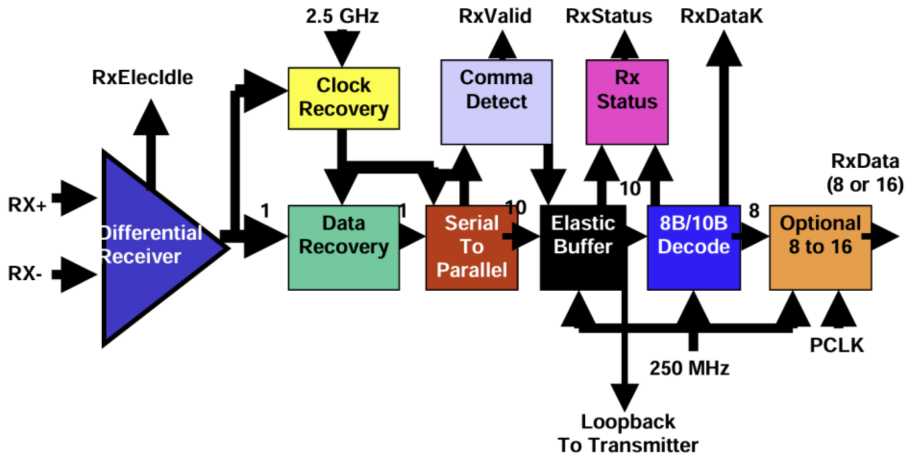

Data Link Layer

- Scope: Logic-blocks, Subsystems

Ensures functionalities across links.

- Functionalities:
  - Error Detection and Correction
  - Framing (Parity Bits)
  - Flow Control
  - Physical Addressing

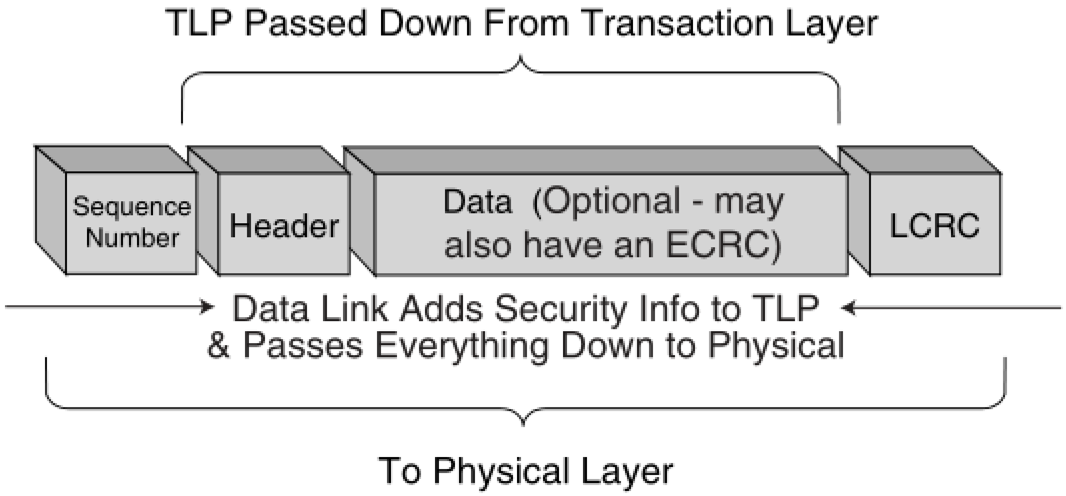

Transport / Transaction Layer

- Scope: End-End, Domain Specific Systems
- Functionalities:
    - Packet Formation
    - Routing
    - Transaction Management

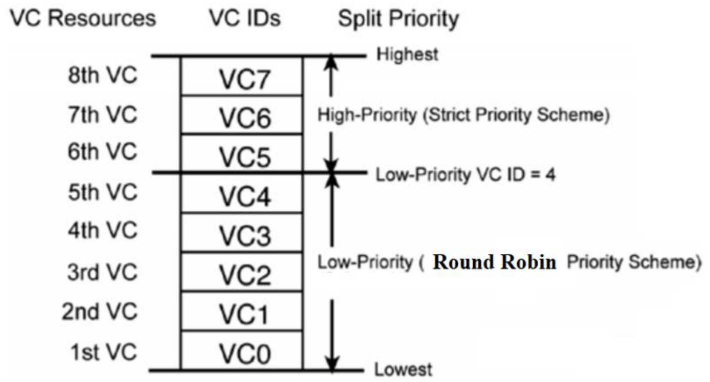

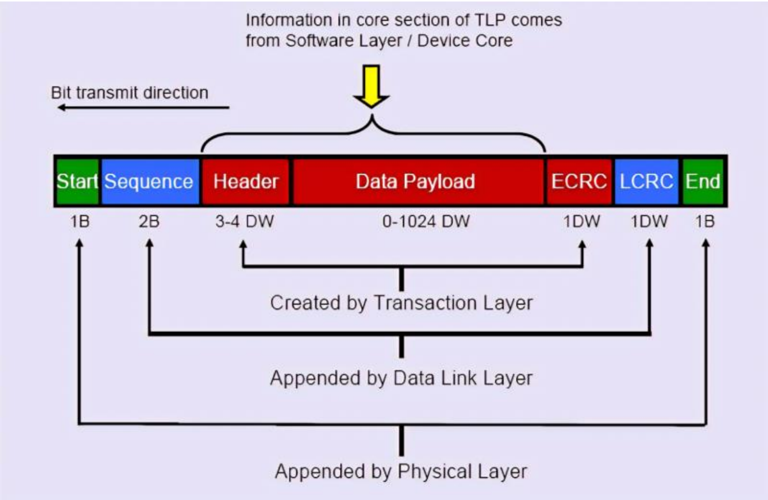

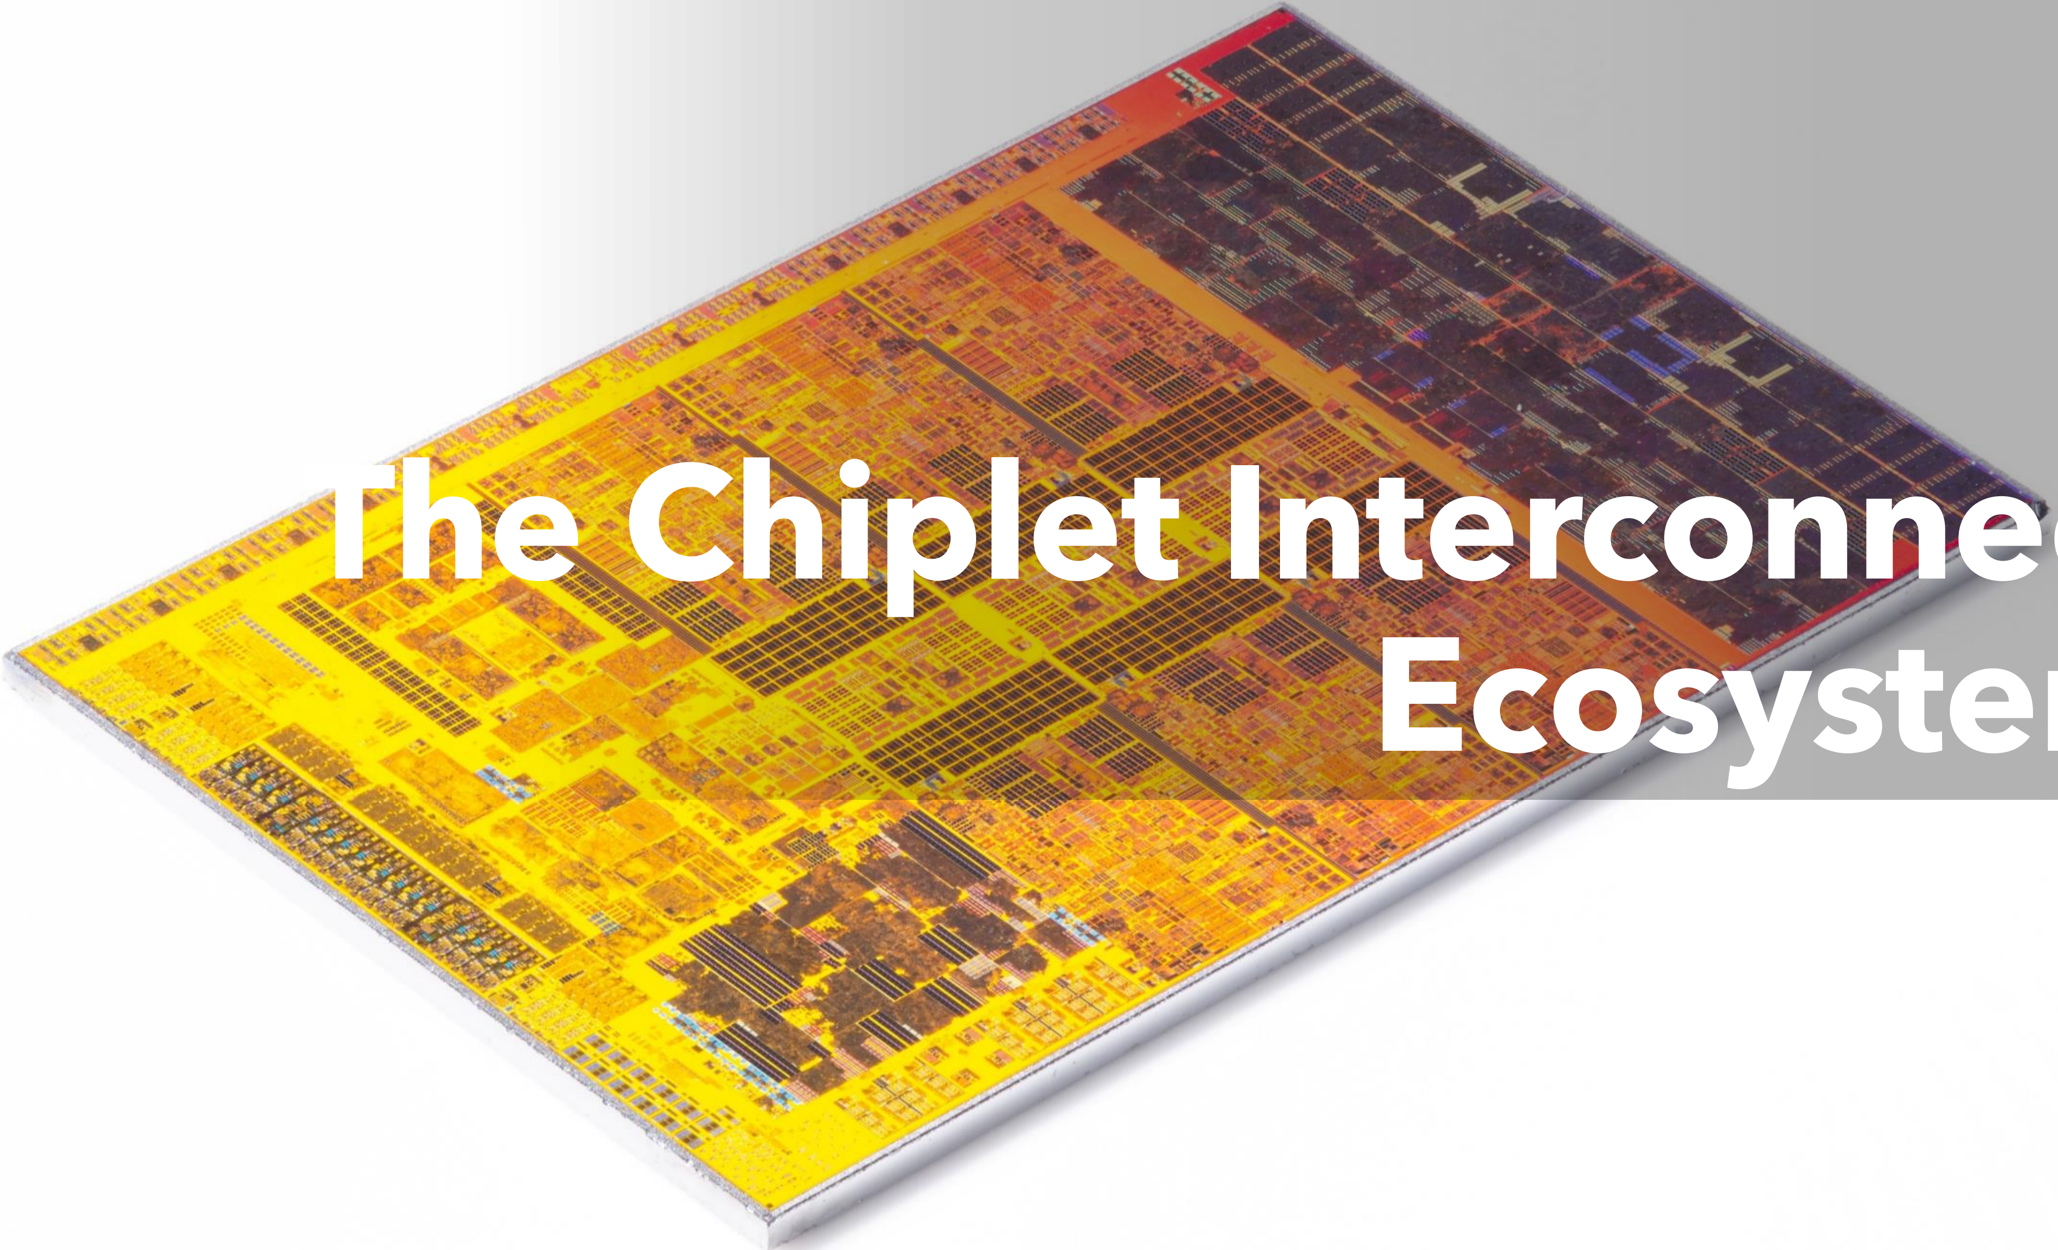

The Chiplet Interconnect Ecosystem

The Great Convergence

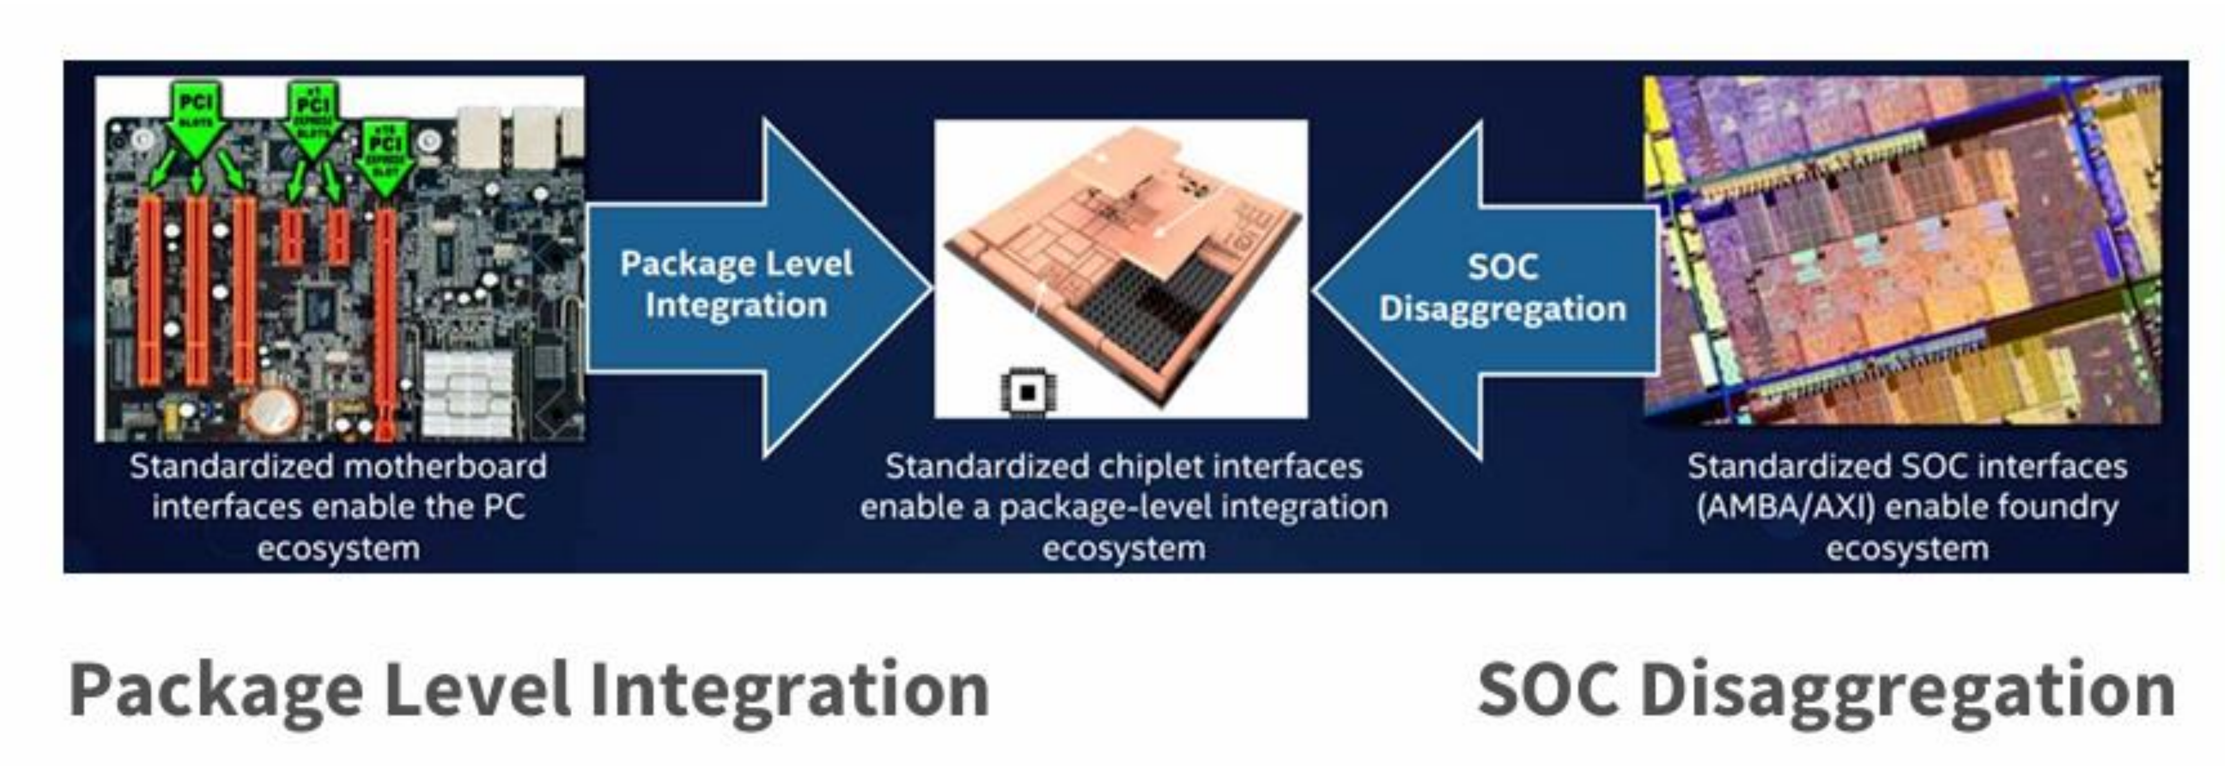

Package Level Integration

SOC Disaggregation

Standardized motherboard interfaces enable the PC ecosystem

Standardized chiplet interfaces enable a package-level integration ecosystem

Standardized SOC interfaces (AMBA/AXI) enable foundry ecosystem

Package Level Integration

SOC Disaggregation

The Biggest Problem

No universal PHY solution to connect chiplets!

We don’t know how to integrate across vendors.

Options:

1. BoW
2. PCIe
3. AIB
4. CNRZ-5

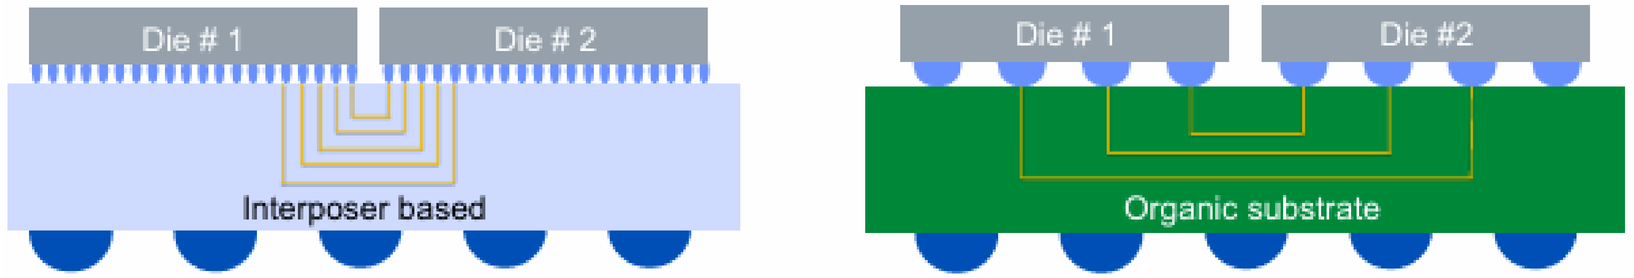

- Higher pin count, parallel bus
- E.g. PHYs: AIB, BoW
- Lower speed bus
- Higher cost packaging

- Lower pin count, serial bus
- E.g. PHYs: AXDieIO, CNRZ-5
- Higher speed bus
- Lower cost packaging

Open Domain-Specific Architecture Project

- Key driving force behind the creation of a "Chiplet Economy"
- Focus on two things:
  - Disaggregation of SoC into SiP
  - Aggregation of a board to a package
- Three pillars:
  - Develop open die to die interfaces
  - Build reference designs
  - Standardize reference workflows

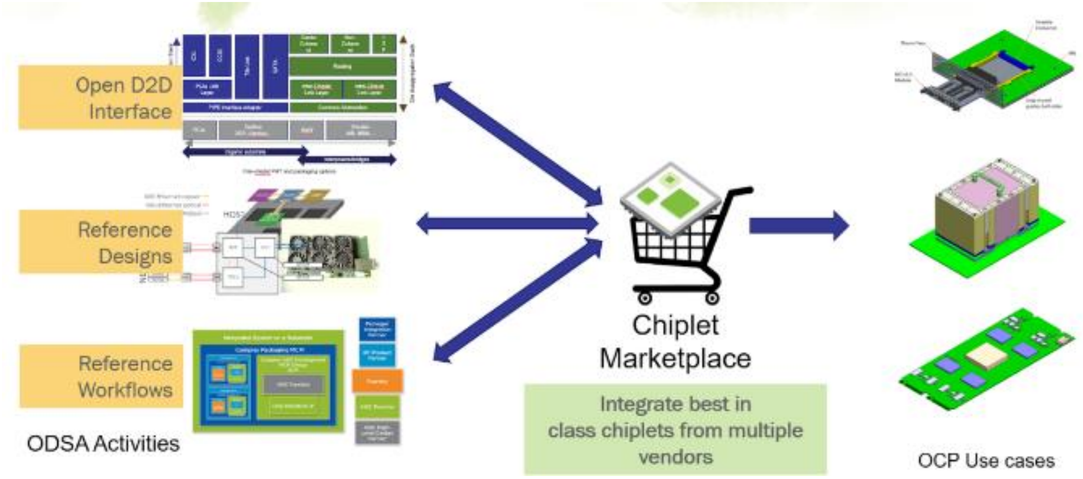

PIPE – Physical Interface for PCI Express

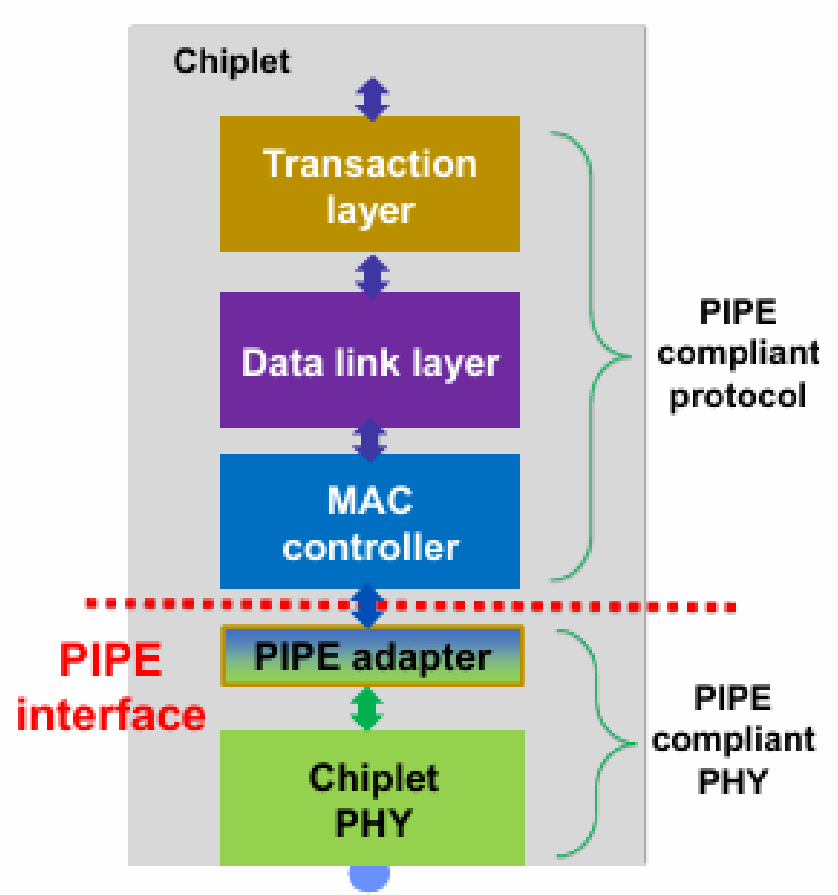

Created by Intel

Provides functional separation between:
1. Logical Layers
2. Electrical interface

Why do we want this?
1. Specialization
2. Modularity
3. Standardization

The Ecosystem and Architecture

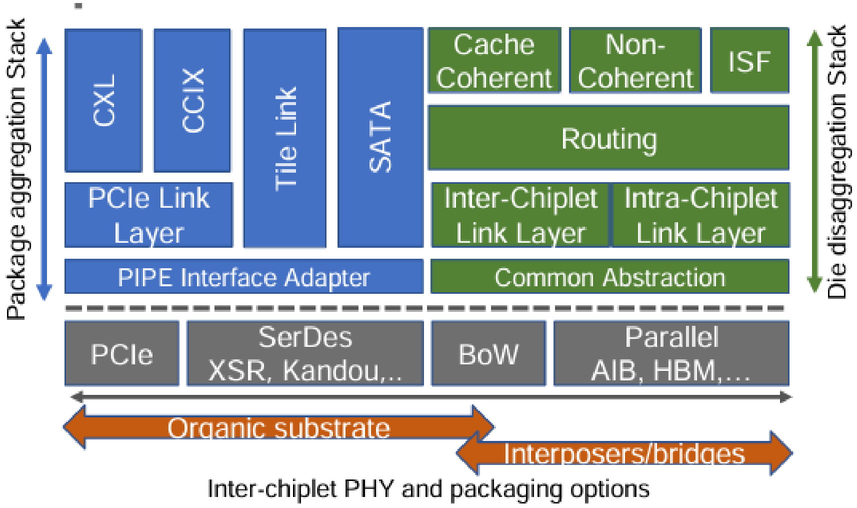

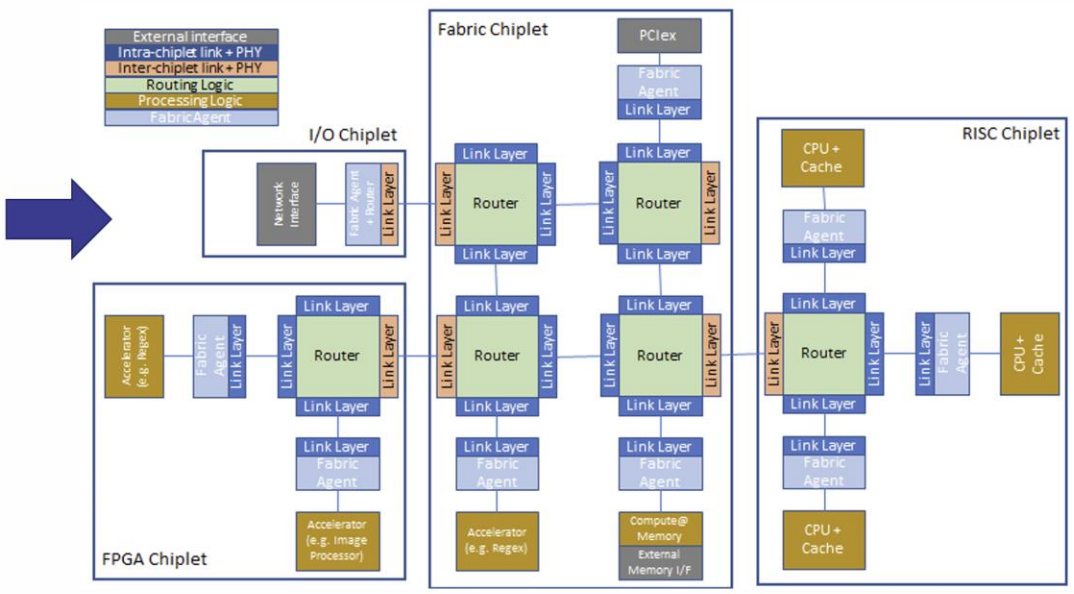

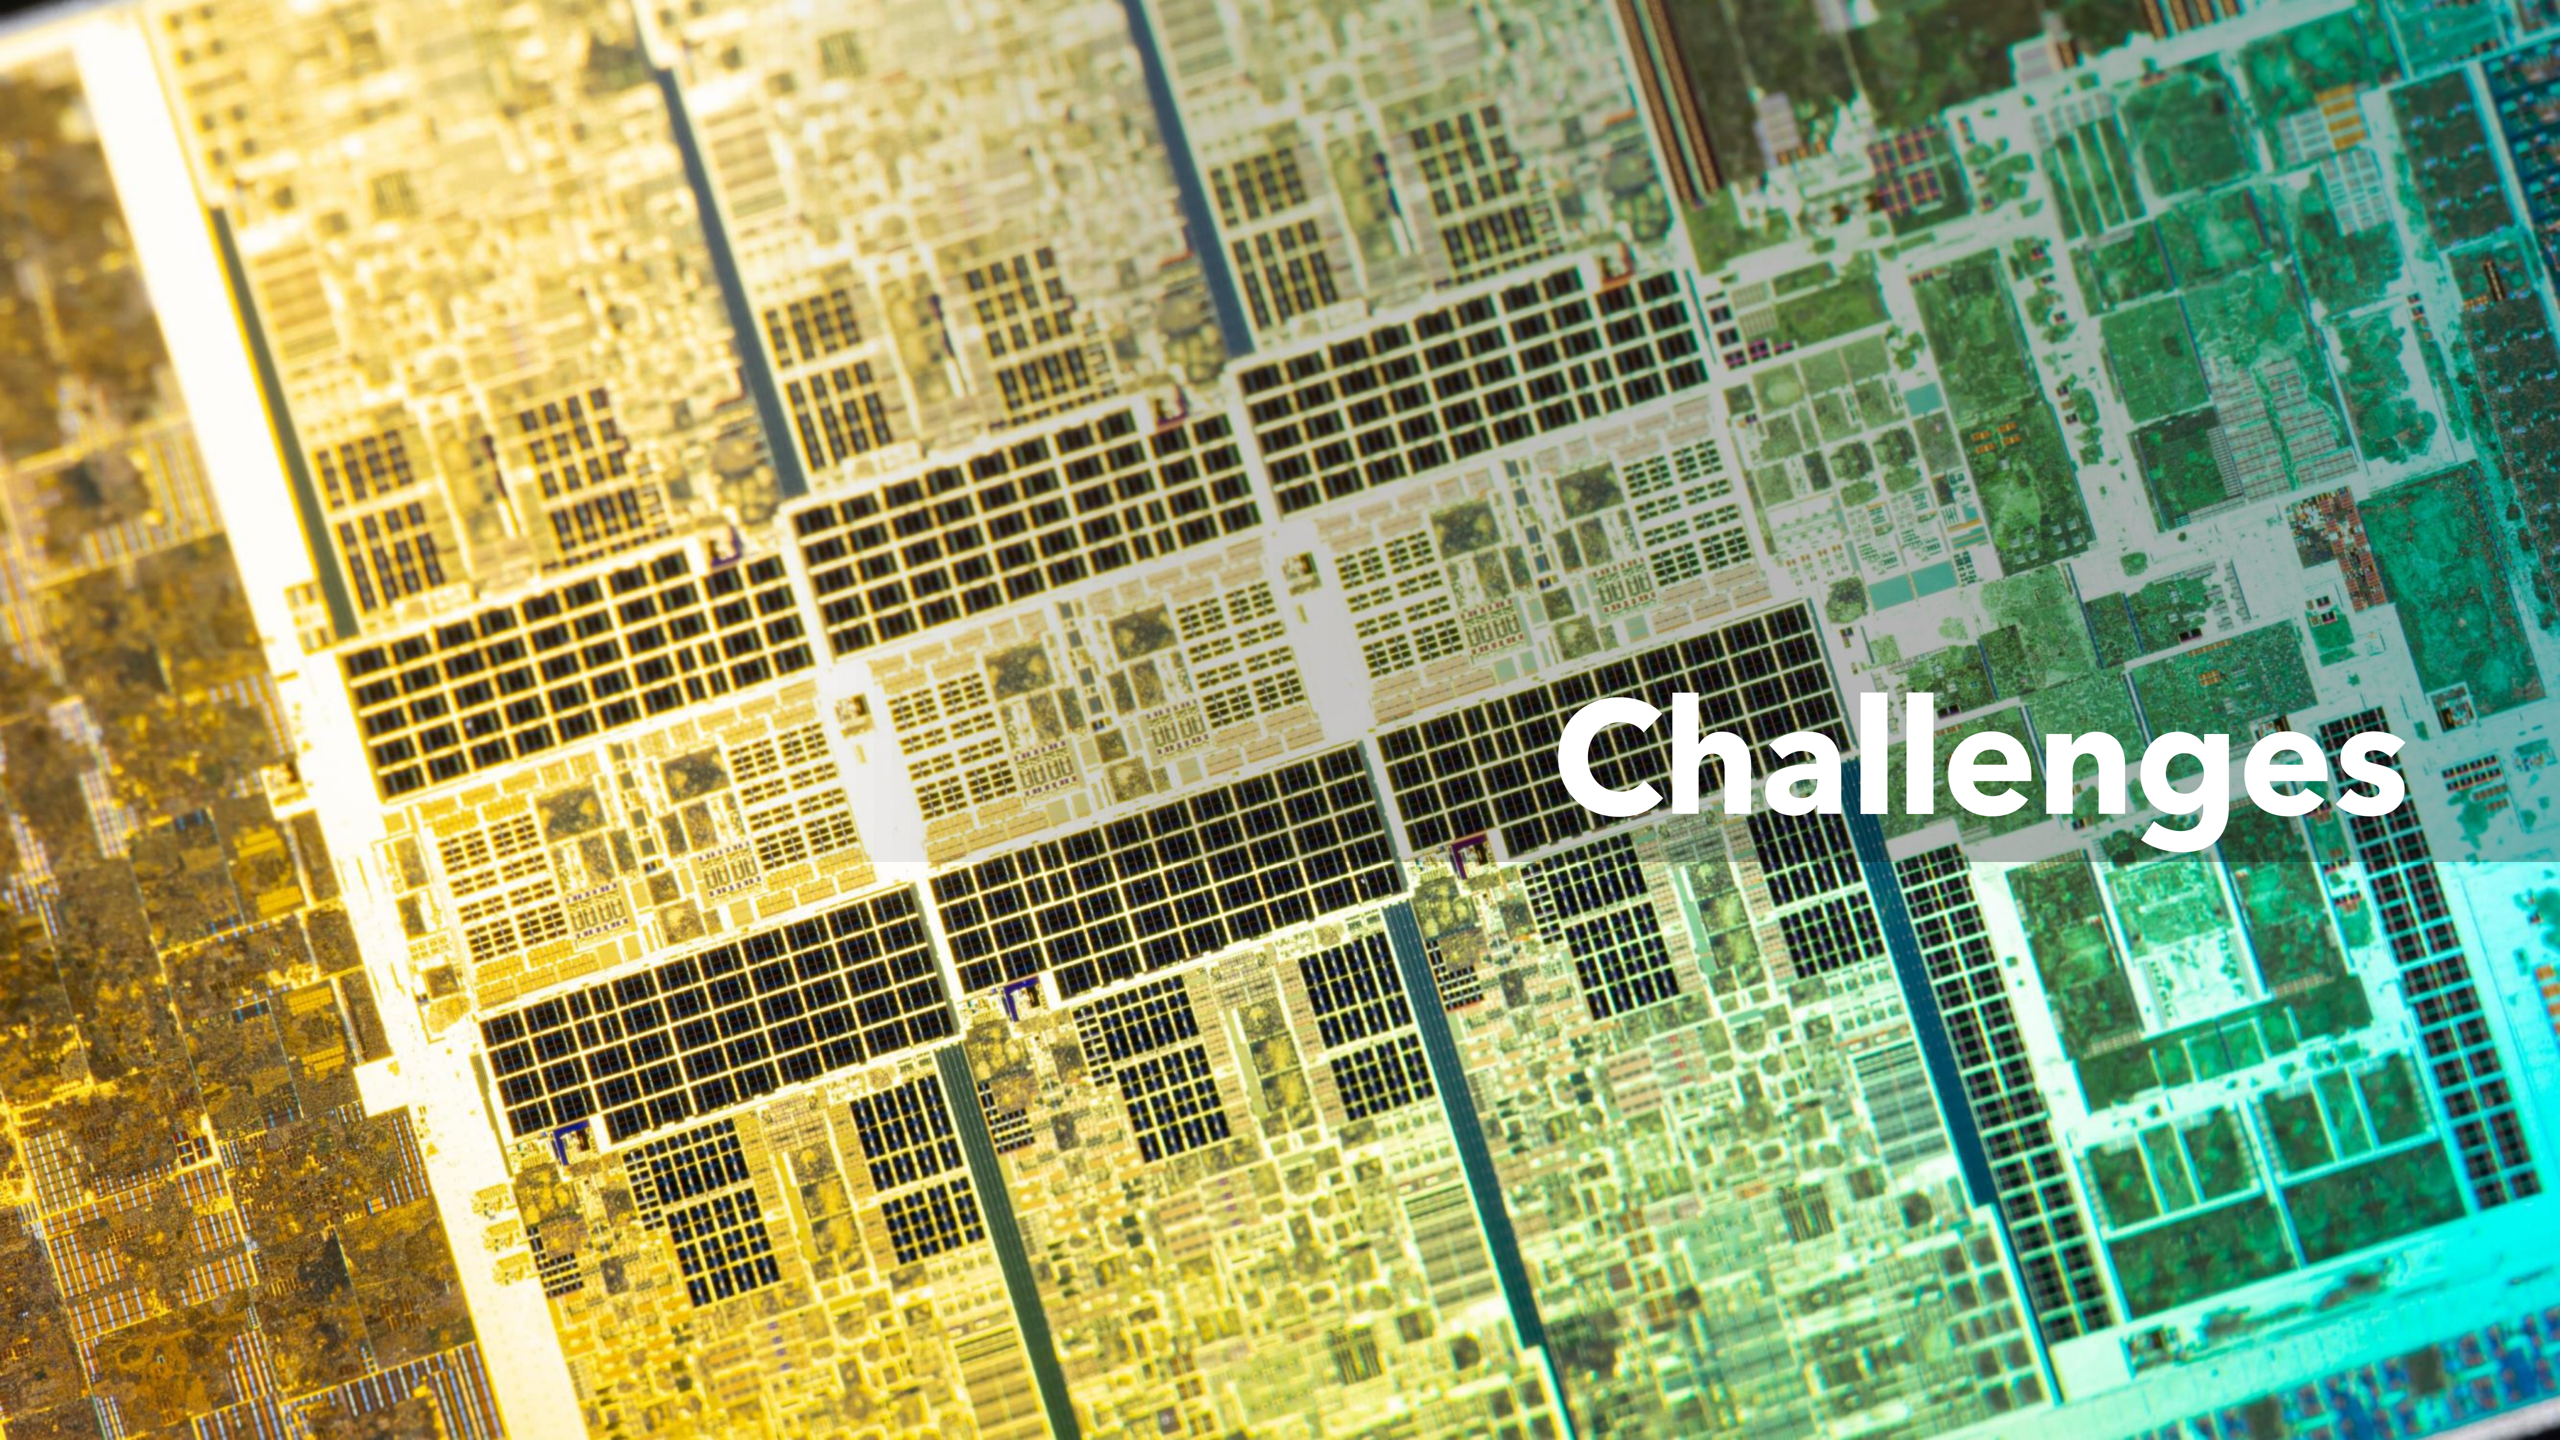

Challenges

Multi-Vendor Collaboration

Simulate

- EDA tool's support for chiplet simulation and emulation
- Data exchange for cross-chiplet simulation

Build

- IP Availability
- Chiplets Inventory
- Chiplet compatibility

Integrate

- Interfaces between chiplets
- Chiplets Yield
- Packaging and substrate compatibility
- Testing and performance

Deliver

- Customer sales and delivery
- Customer Support
- RMA and Failure Analysis

Cost

Heterogeneously integrated chips only pay off:

- $ of die defects > $ of packaging

Increasing granularity too much has marginal utility.

Design and integration effort costs more then SoC equivelant.

Advanced packaging methods struggle with yield.

Co-Design

Vendors may not be willing to share intellectual property

Can you design and integrate with black boxes?

Testing

We cannot detach a bonded die from a substrate or interposer after packaging

Chips have to be independently verified

Chips may have "correctness" issues

Security

Larger Attack Surface

Untrusted Supply Chains

Difficulty adopting Standard Protection Mechanisms

That is all! Join Us!

Have an Asianometry video

[Why AMD's Chiplets Work (youtube.com)]