Road to Chiplets: Architecture

July 13 & 14, 2021

MEPTEC
THE NEXT GENERATION

www.meptec.org

Pathfinding and Design of Large-Scale Chiplet-Based Systems

PUNEET GUPTA (PUNEETG@UCLA.EDU)
UNIVERSITY OF CALIFORNIA, LOS ANGELES

How much can you gain by the Chiplet approach?

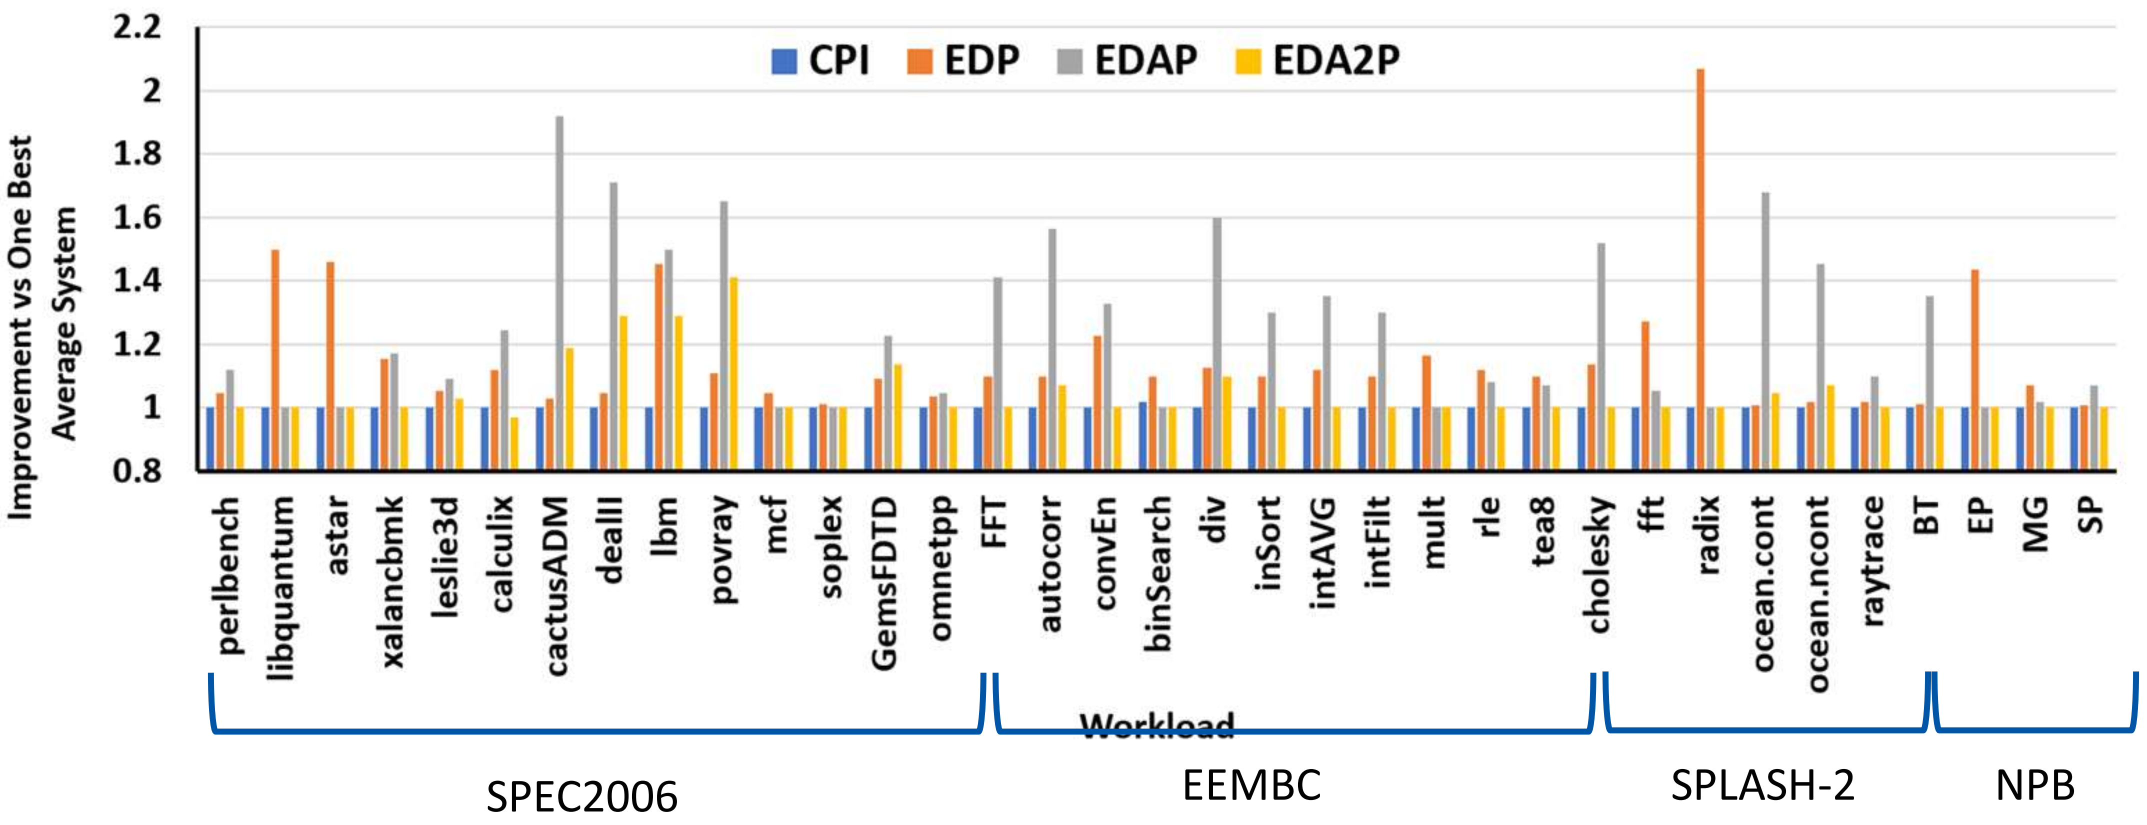

Up to 2.1X gain by system customization!

- Hypothetical study based on processor + cache chiplets and application specific customization

Results: SoC vs Chiplet Assembly

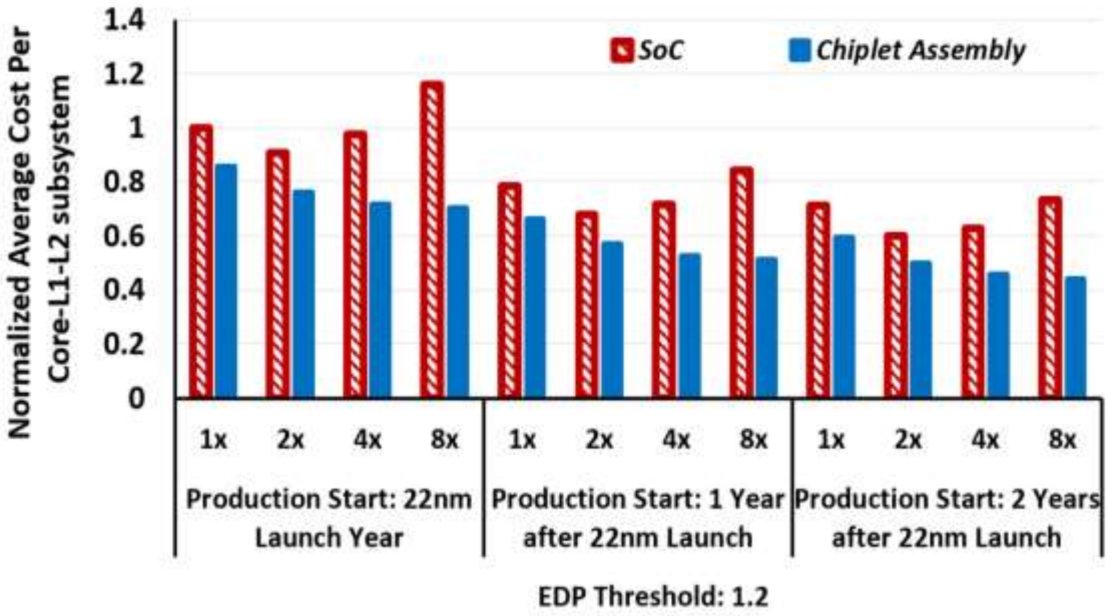

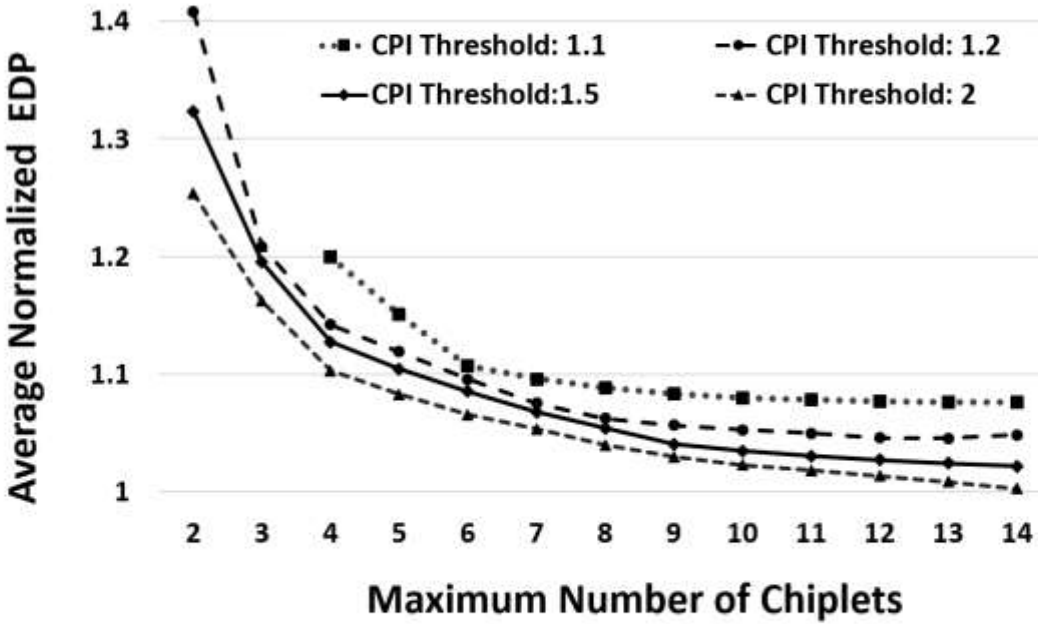

Key observations

- Large multi-core processors if built using chiplets can provide significant cost savings
  - Small chips don’t benefit
  - What qualifies as “large” depends on technology maturity and integration costs
- Overall cost benefit drops if you are not willing to take a (small) performance hit
  - “Ecosystem” cost depends also on reuse across multiple products

Designing Large Chiplet-based Systems

TO APPEAR IN DAC’21

Waferscale Integration: The Chiplet Approach

- High Bandwidth and Energy-efficient Communication:
  High density interconnection

- Large amount of Computation as well as Memory Capacity
  Heterogeneous integration of compute and high-density memory

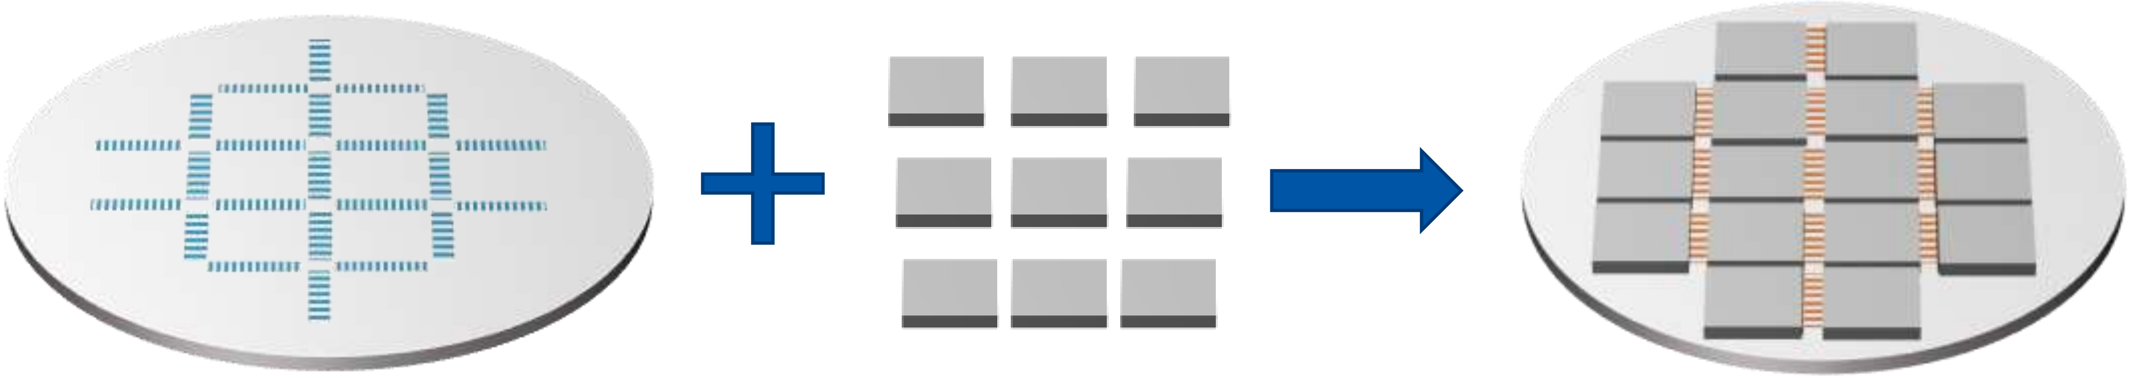

A wafer with interconnect wiring only

Small known good chiplets

Bond the dies on to the interconnect wafer

Enabling WSI Technology

UCLA Silicon Interconnect Fabric (Si-IF)*

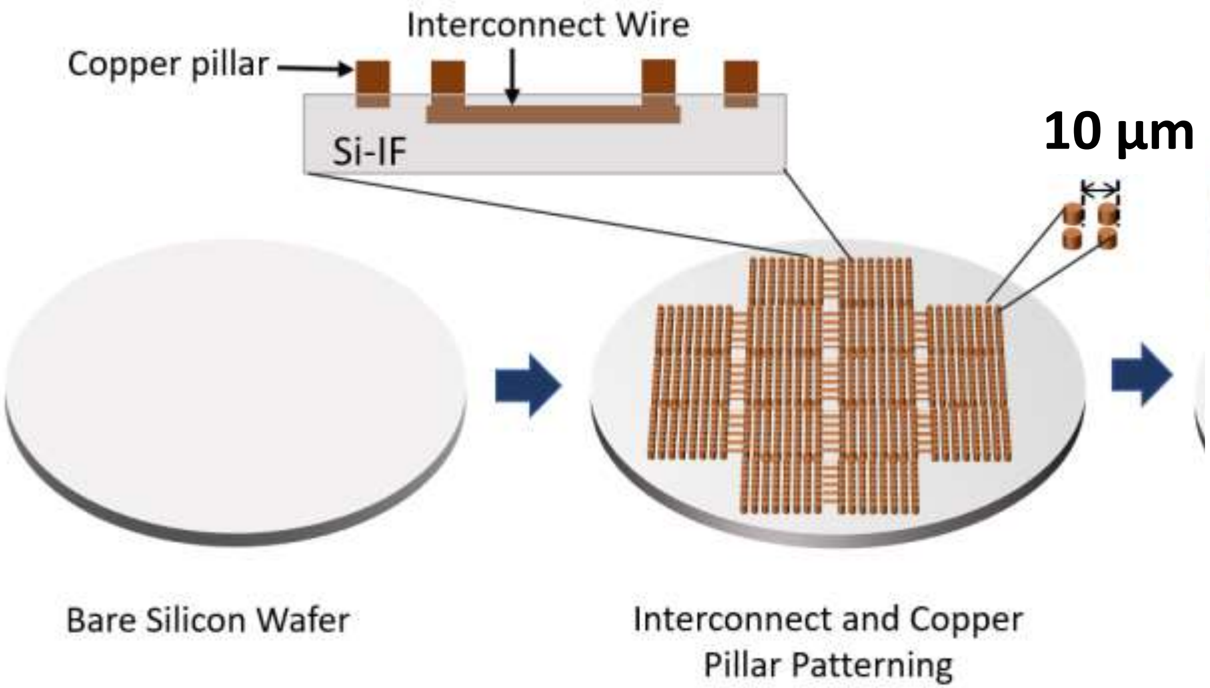

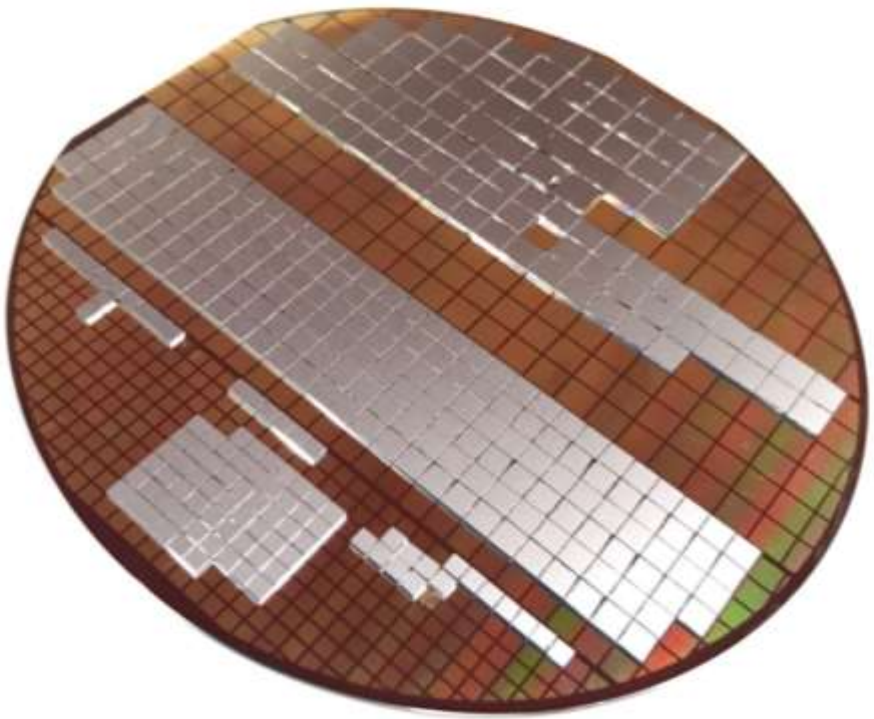

Measured Bond Yield >99%

Allows waferscale integration with high yield

*UCLA CHIPS Program: https://www.chips.ucla.edu/research/project/4

2048 Chiplet Architecture

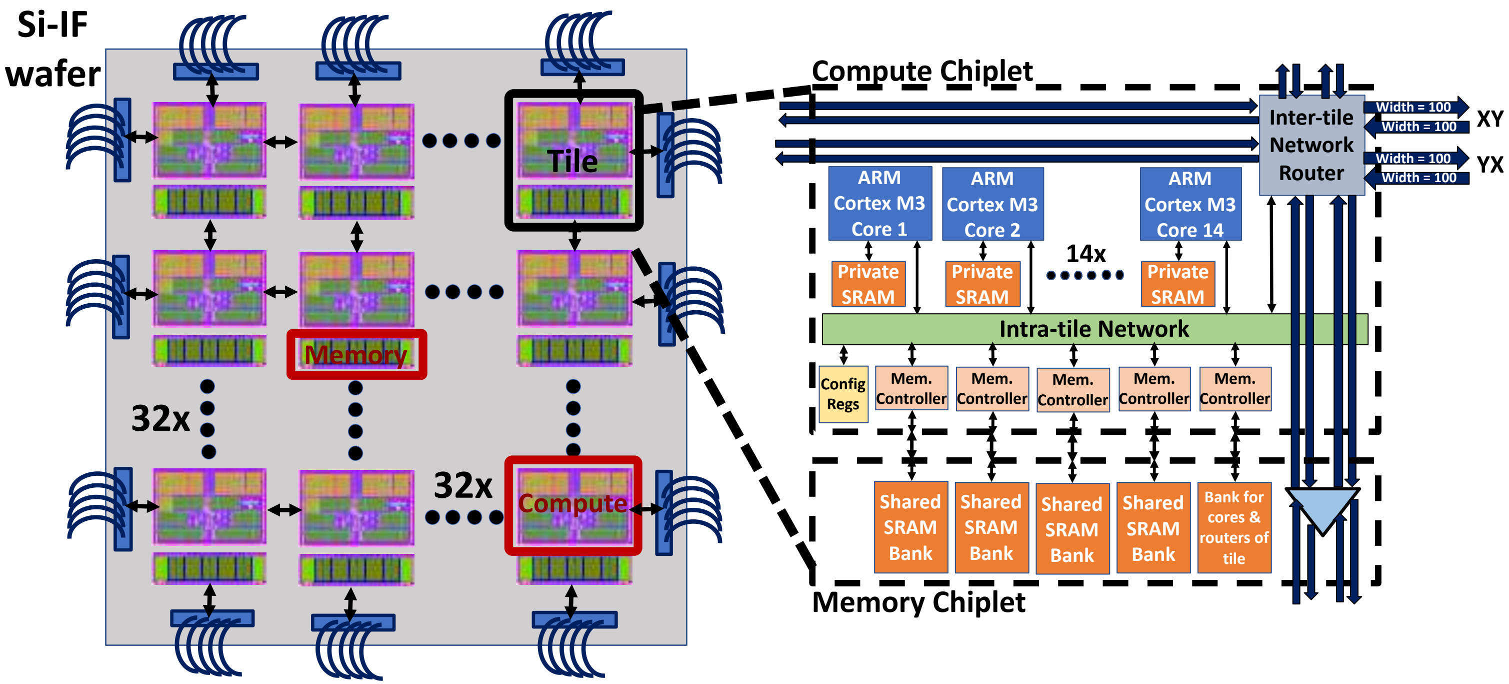

- Implemented in **TSMC N40-LP**
- Tiles : **1024 (Total 14,336 Cores)**
- Private memory per Core: **64KB**
- Total Shared Memory: **512MB**
- Shared Memory Bandwidth : **6.14 TB/s**
- Network Bandwidth: **9.83 TB/s**
- Total Compute : **4.3 TOPs**
- Peak Power : **725W**
- Total Area: **15100 mm^2**

Challenges Faced While Designing the System

How should we **deliver power** to all the flip-chip bonded chiplets across the wafer?

How can we reliably **distribute clock** across such a large area?

What is the **testing strategy** for such a large system?

What is the **inter-chip network architecture** and how do we achieve resiliency if a few chiplets fail?

Power Delivery

- Edge Power Delivery at 2.5V
  - Wasteful but simple
  - Back or front side power delivery possible but more complex
- LDO based power management at each node
- On-chip decoupling capacitance (20nF per tile)
- DeCap consumes 30% of the chip area
  - *Deep Trench Capacitors would help*

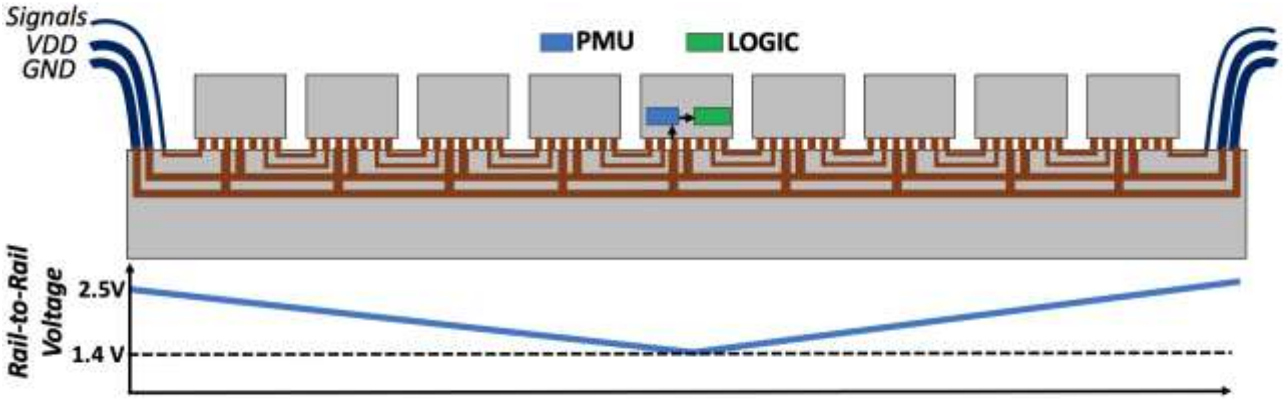

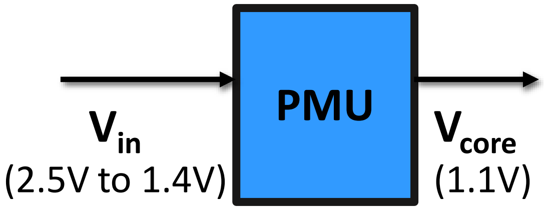

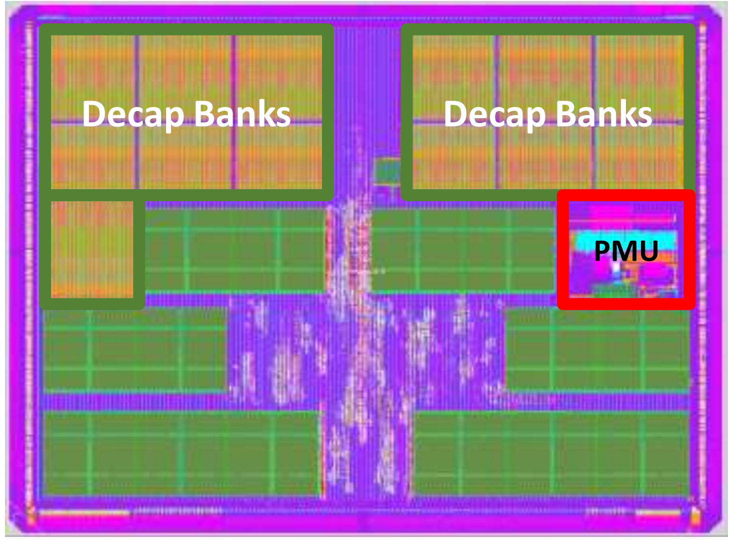

Waferscale Clocking

- Clock generation
  - Stable reference voltage needed by PLL not present away from edge
  - Generate fast clock at the edge and distribute
- Clock distribution
  - Fast clock is forwarded
  - Clock inverted at each hop to avoid duty cycle distortion accumulation

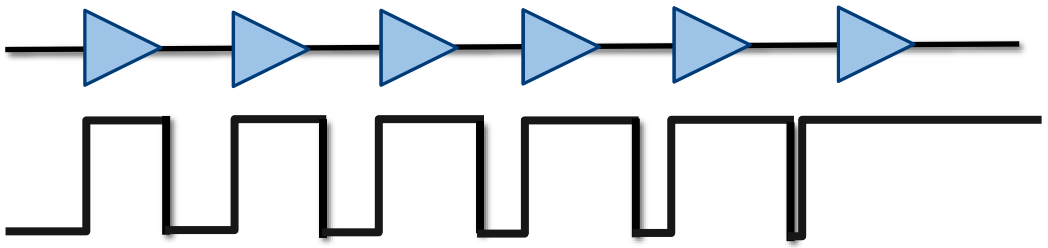

  - Communication between dies using asynchronous interfaces
  - Fault tolerance in clock distribution network

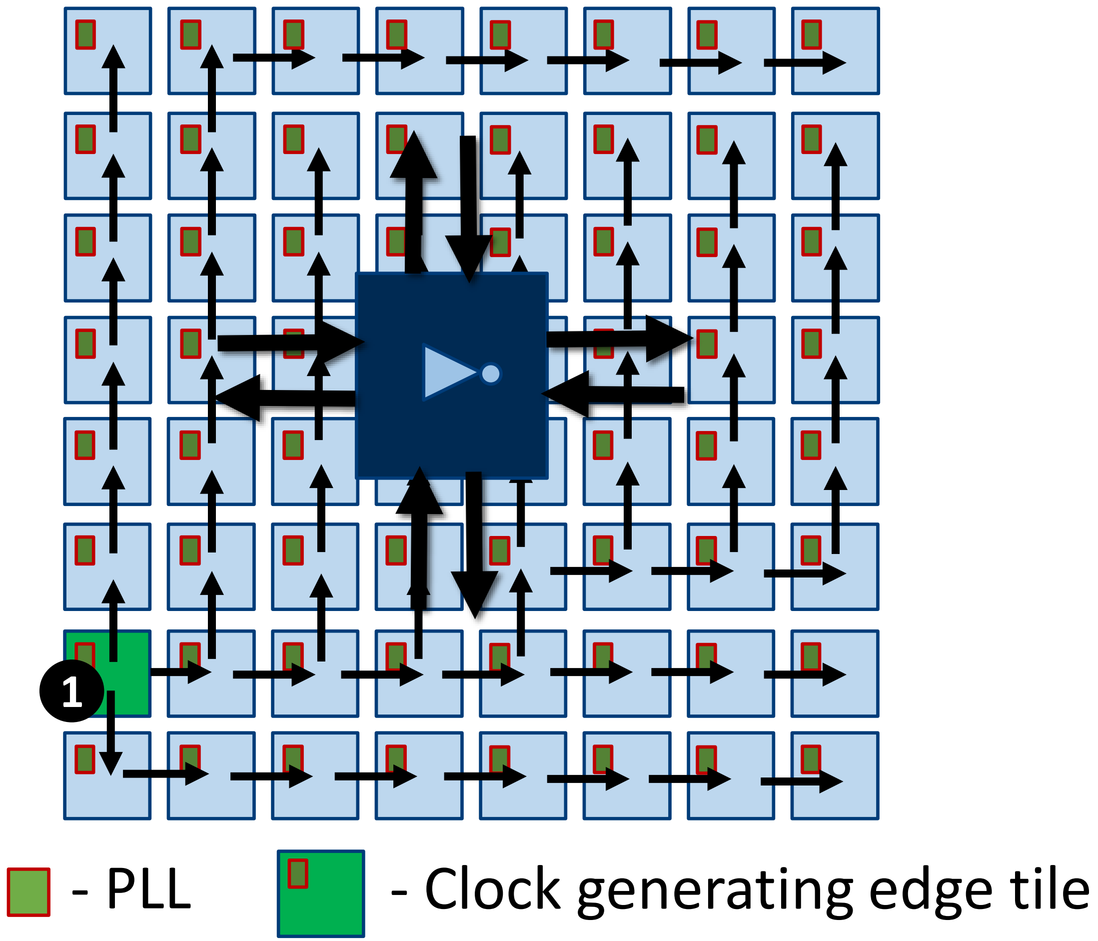

Pre-bond Die Testing

- Fine pitch pads cannot be probed
- Larger sacrificial pads for probe test

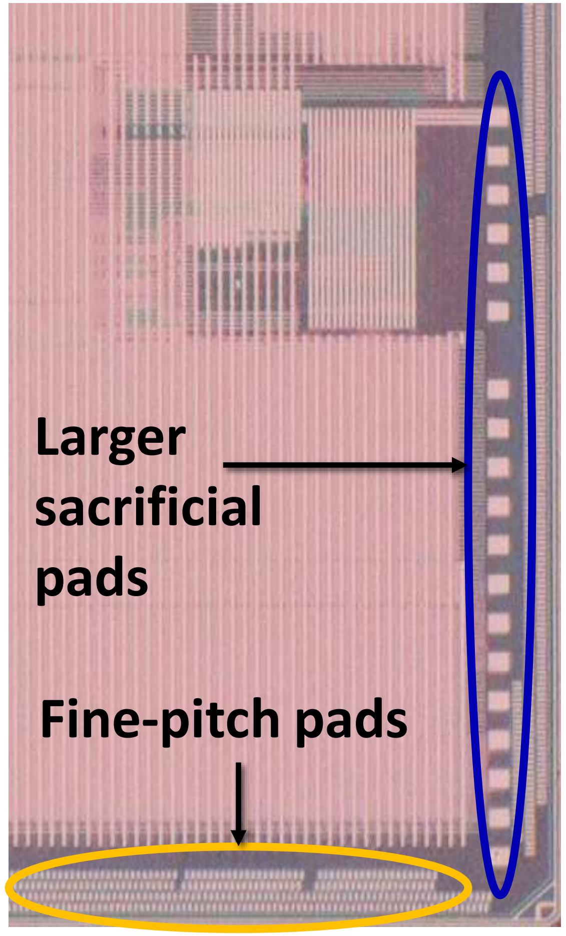

Post-bonding JTAG Test Scheme

- (1) Multiple chains
  - One JTAG chain results in single point of failure vulnerability
  - Throughput is an issue:
    - 2.5 hours to load the memories using one chain
    - 5 minutes to load with 32 chains

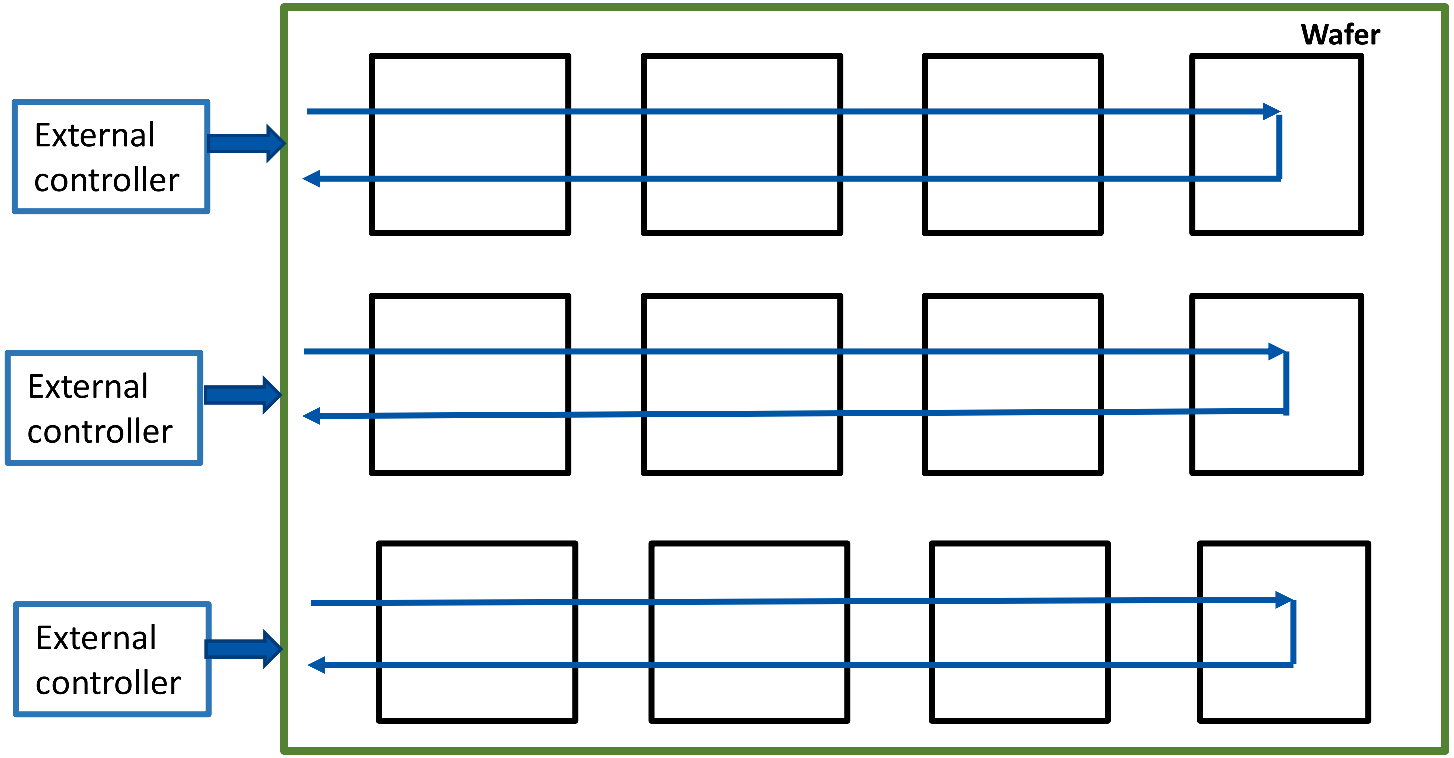

Post-bonding JTAG Test Scheme

- (1) Multiple chains
  - One JTAG chain results in single point of failure vulnerability
  - Throughput is an issue:
    - 2.5 hours to load the memories using one chain
    - 5 minutes to load with 32 chains
- (2) Progressive unrolling
  - Helps identify post-bonding faulty dies

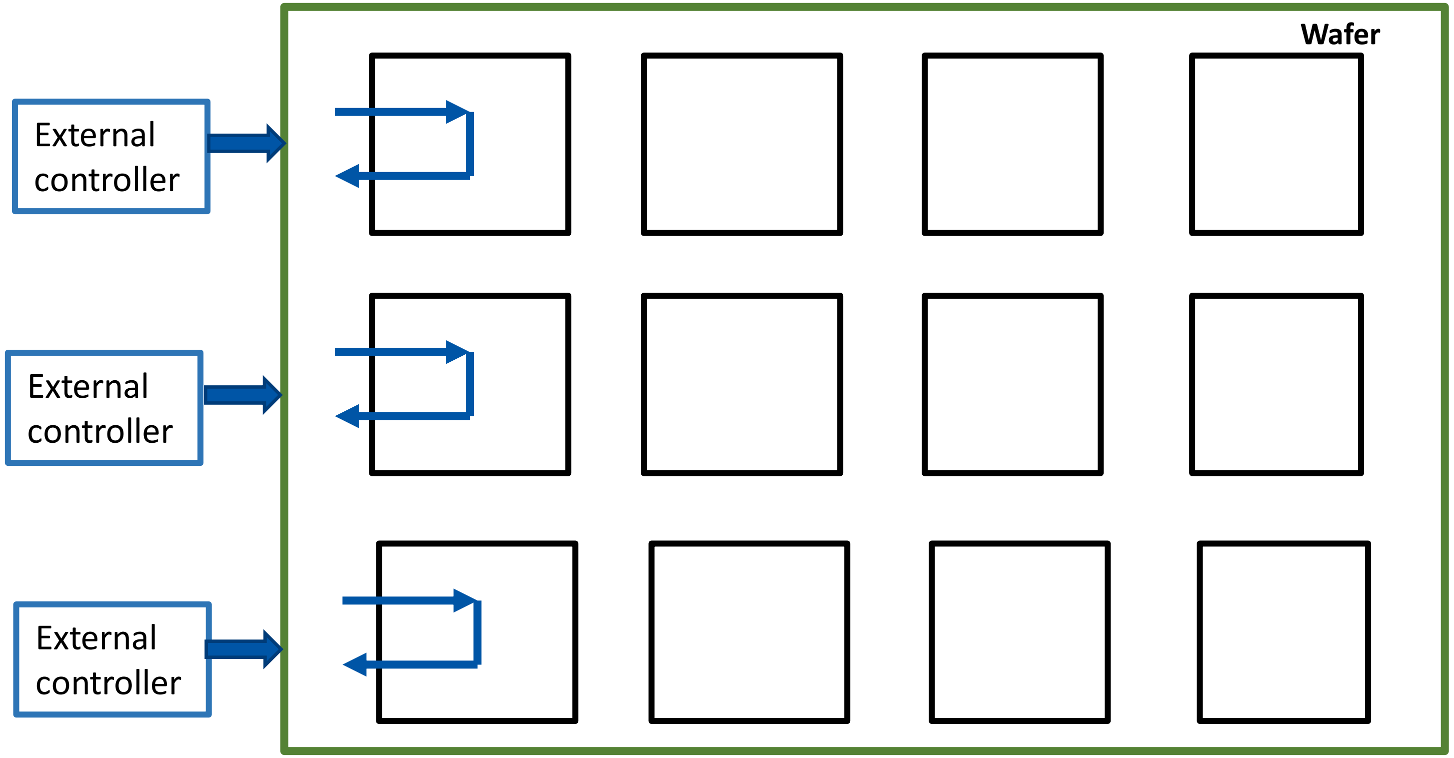

Post-bonding JTAG Test Scheme

- (1) Multiple chains
  - One JTAG chain results in single point of failure vulnerability
  - Throughput is an issue:
    - 2.5 hours to load the memories using one chain
    - 5 minutes to load with 32 chains
- (2) Progressive unrolling
  - Helps identify post-bonding faulty dies

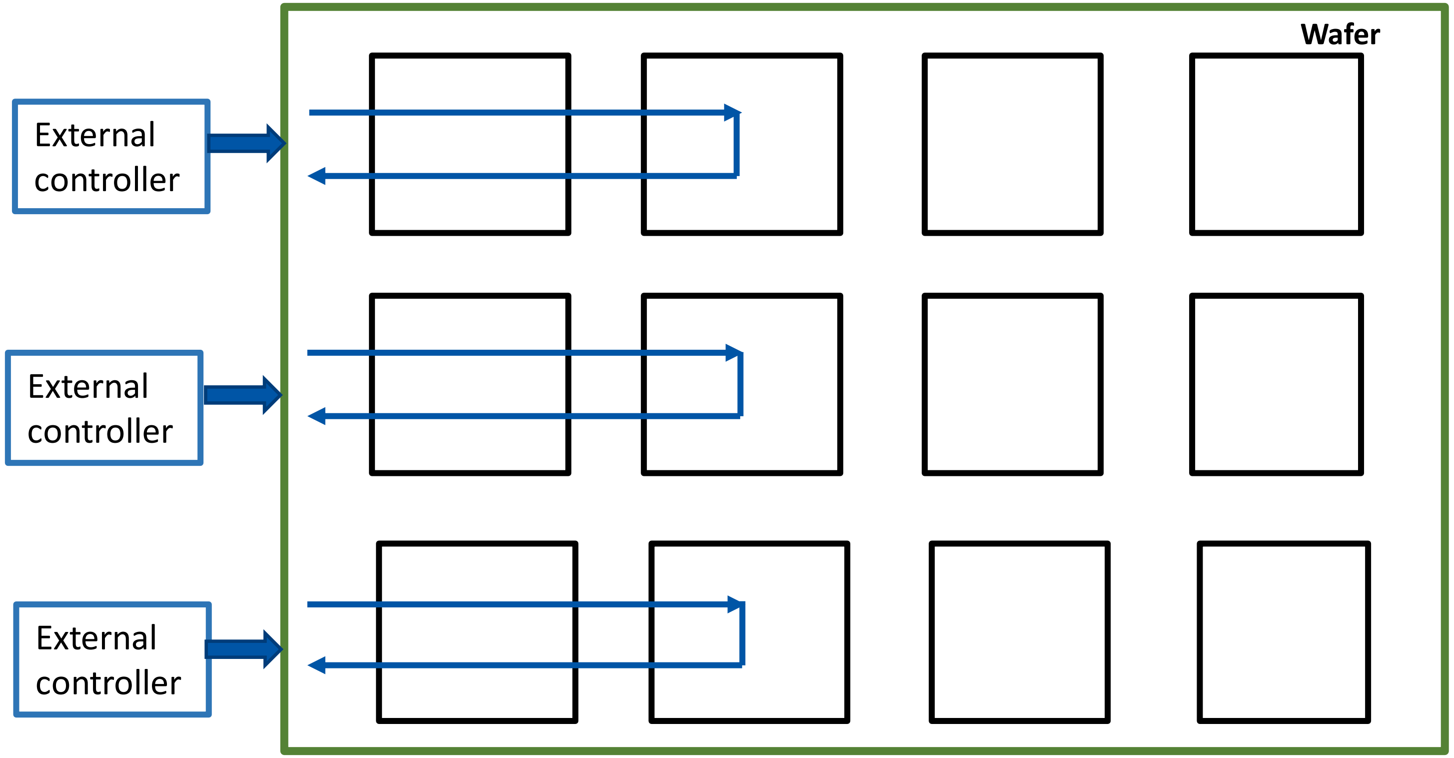

Network Resiliency

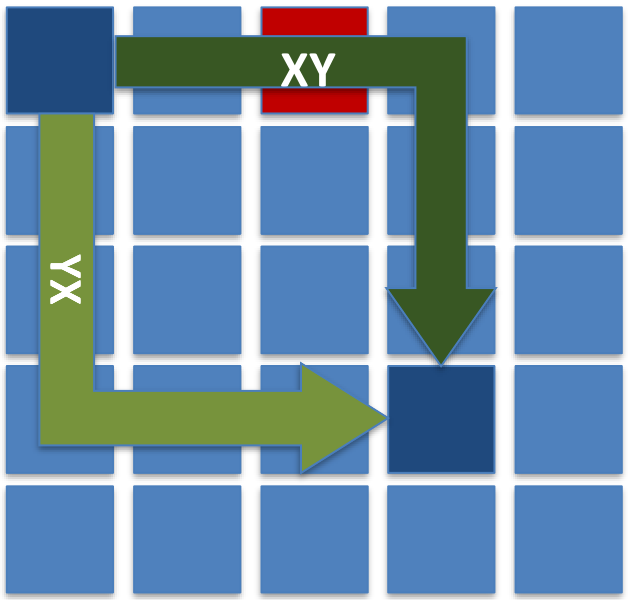

Two Separate Networks

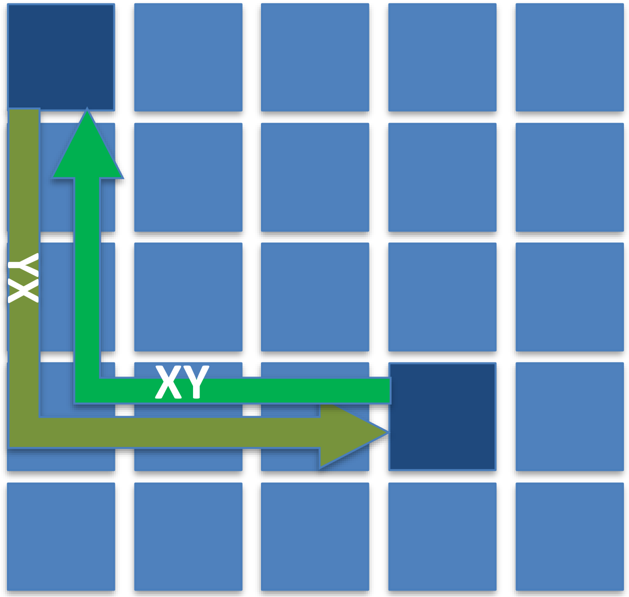

Request-Response in Complimentary Networks

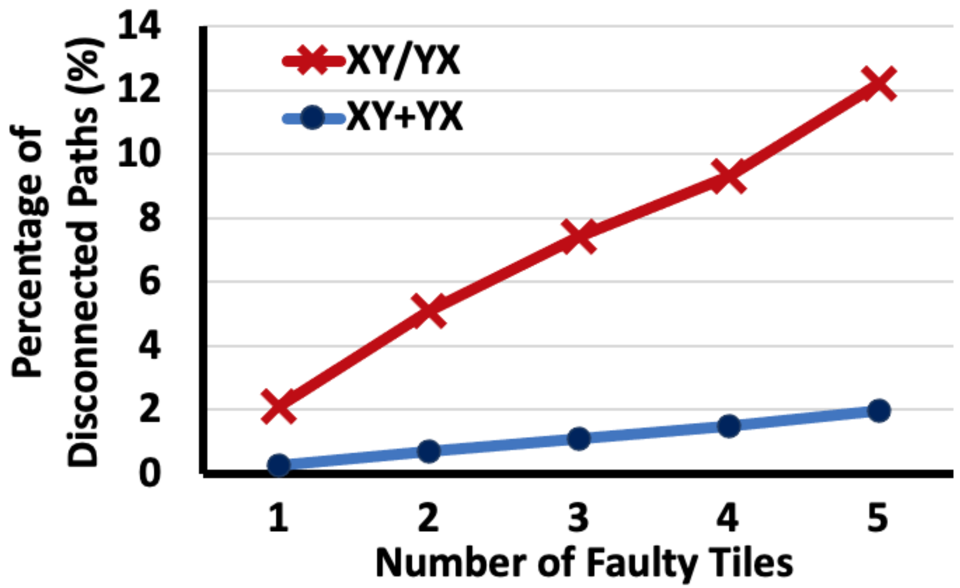

I/O Architecture

- I/O pitch of 10 um and depth of 20 um
- Simple cascaded buffer architecture
- 0.07 - 0.18 pJ/bit
- Two pillars per IO for redundancy
- ESD diodes and buffers need to fit within the I/O footprint

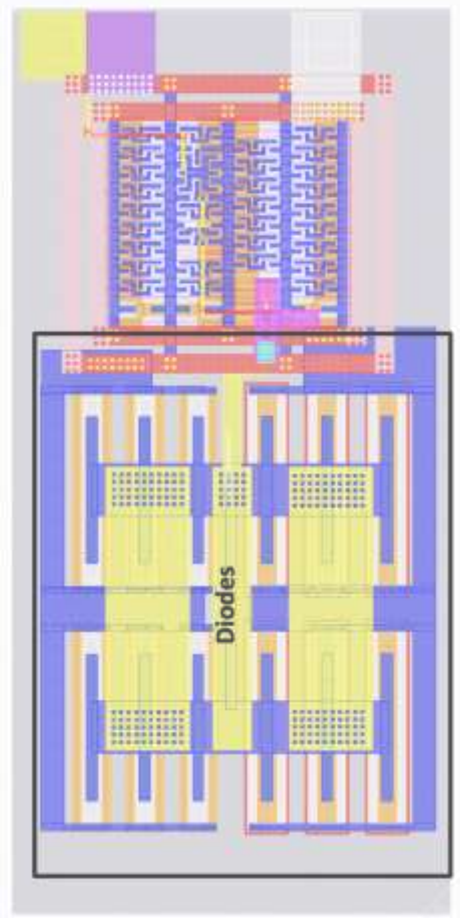

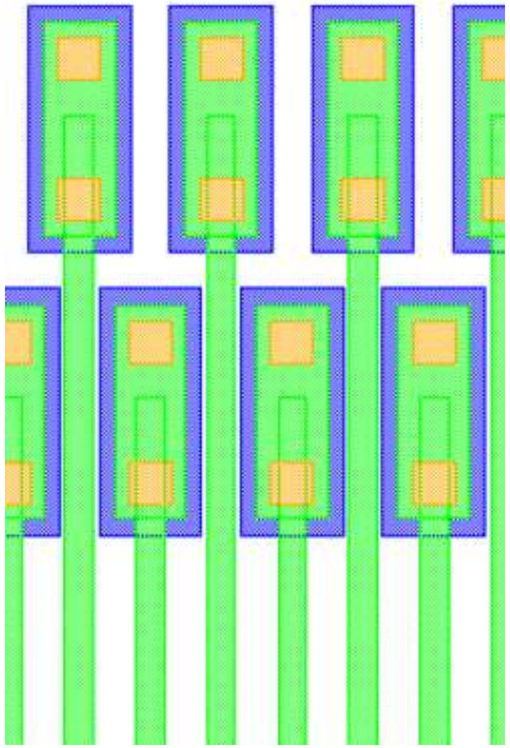

Chiplets Tested, Waferscale System Assembly in Progress

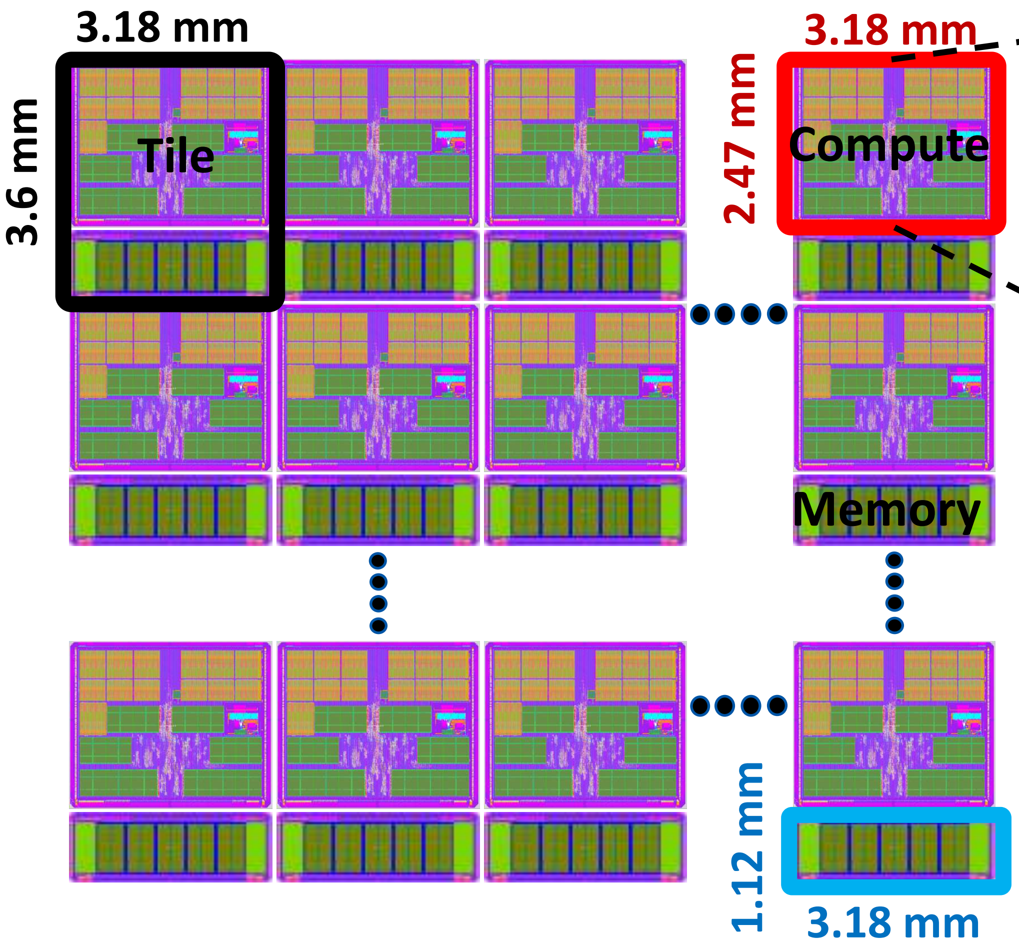

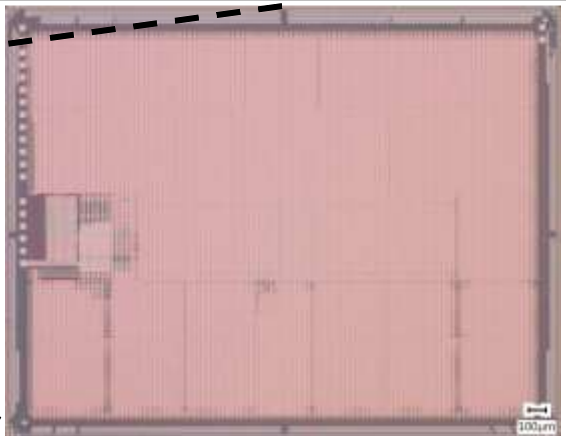

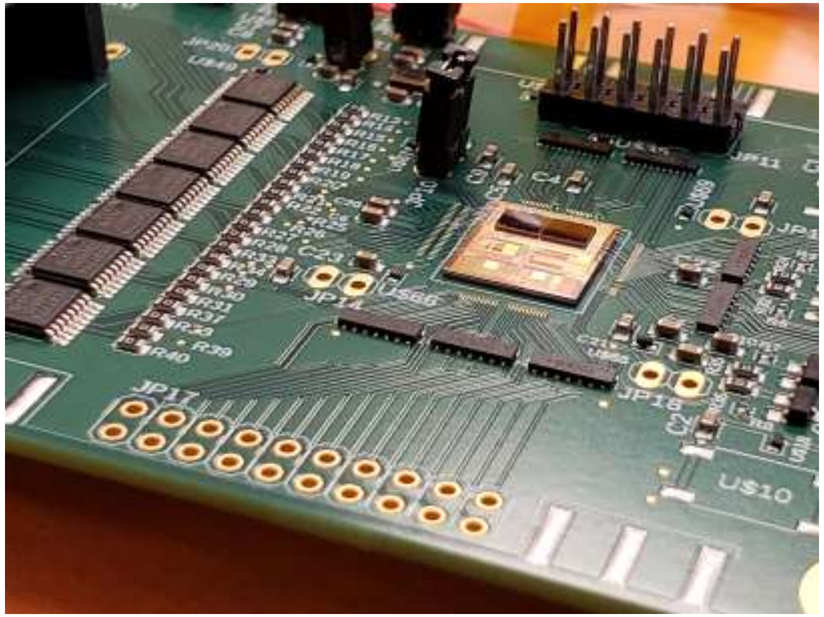

Pathfinding 2.5D Substrates

APPEARED IN SLIP'20

Diverse 2.5D Integration Technologies

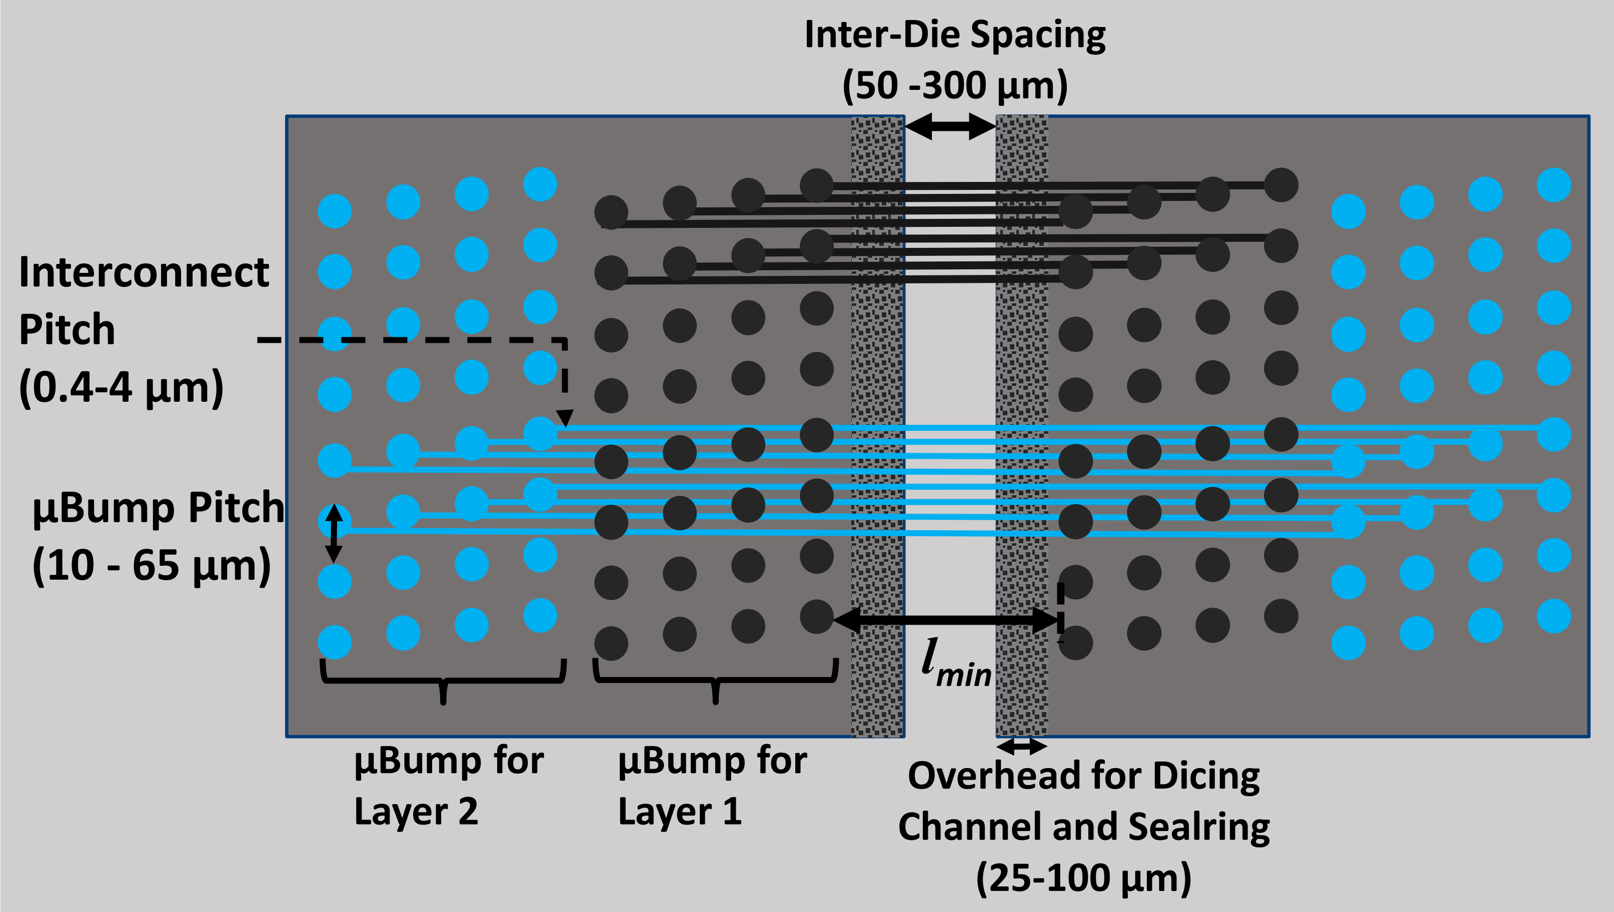

Interconnect Link is Not Just the Wire

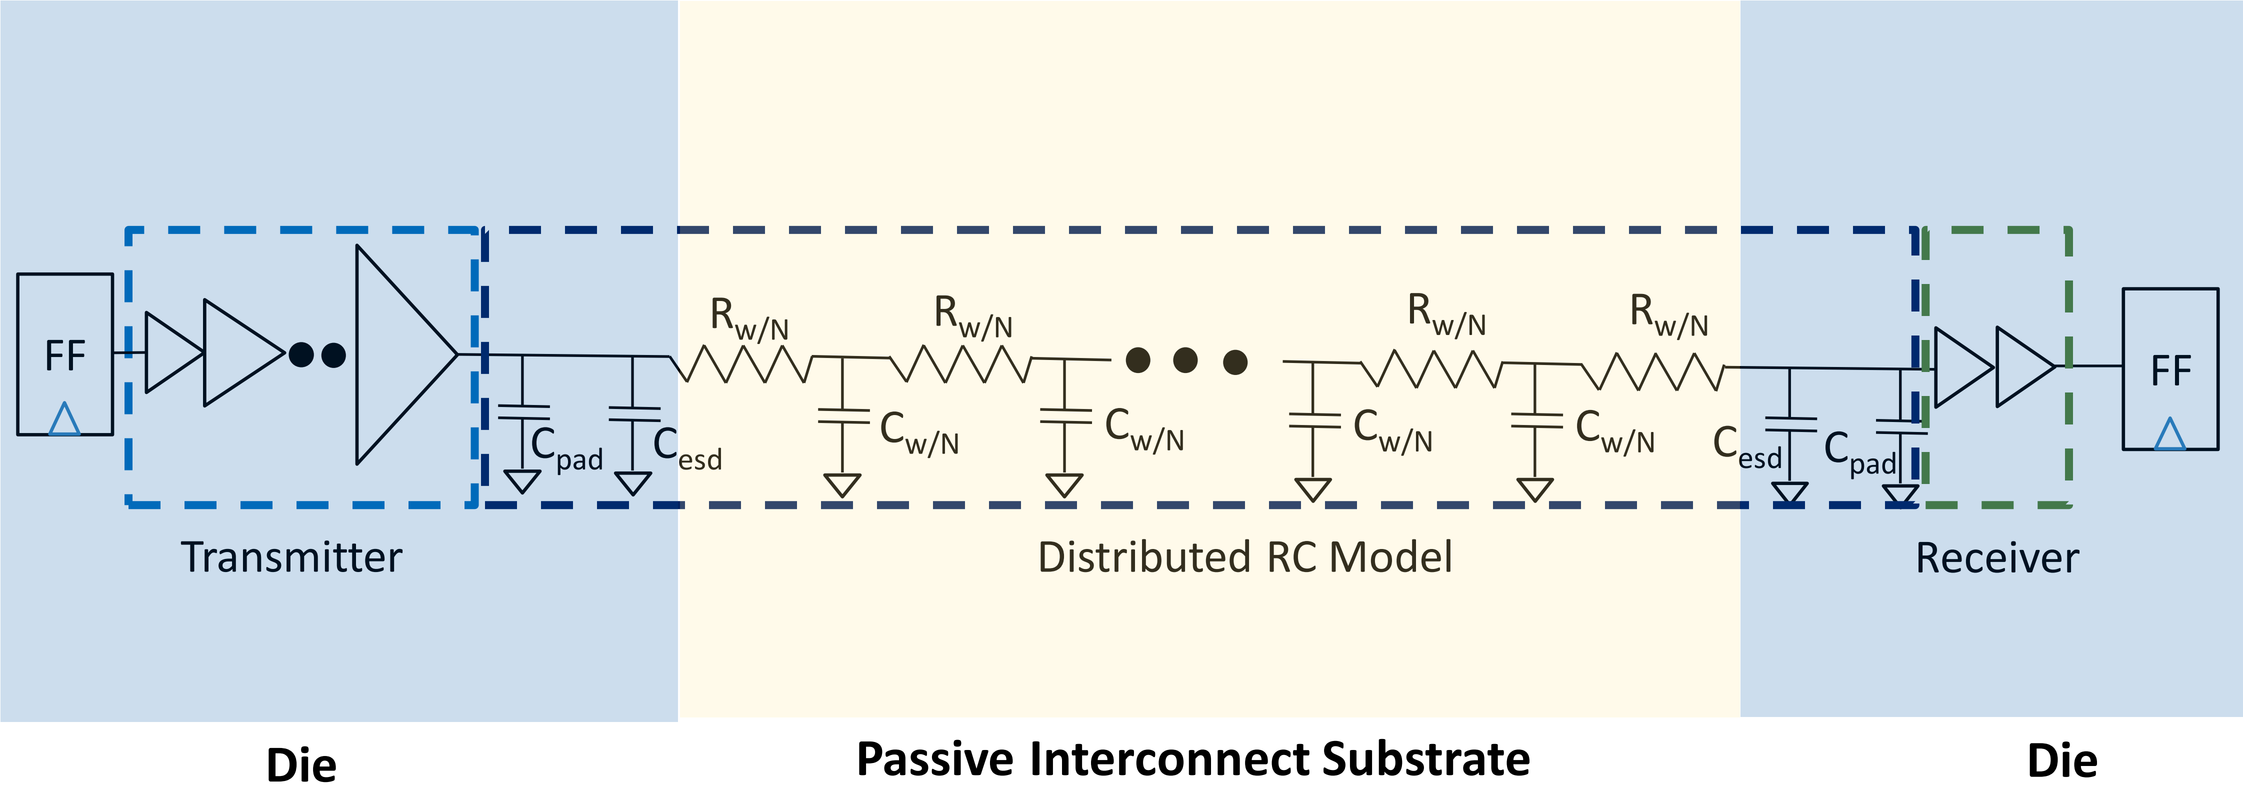

Scaling **µBump pitch** vs **Wire pitch**

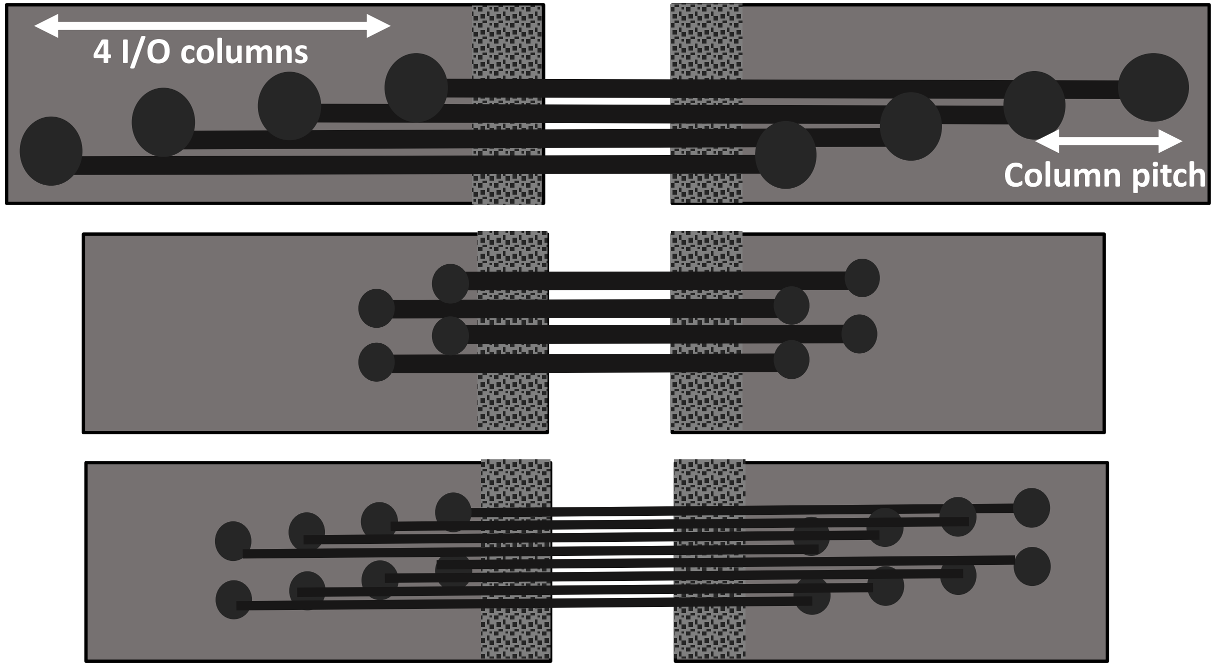

- ***Scaling down the µBump pitch reduces wire-length***
- ***Scaling down the wire-pitch increases wire-length (to maximally utilize wire density)***

Scaling **µBump pitch** vs **Wire pitch**

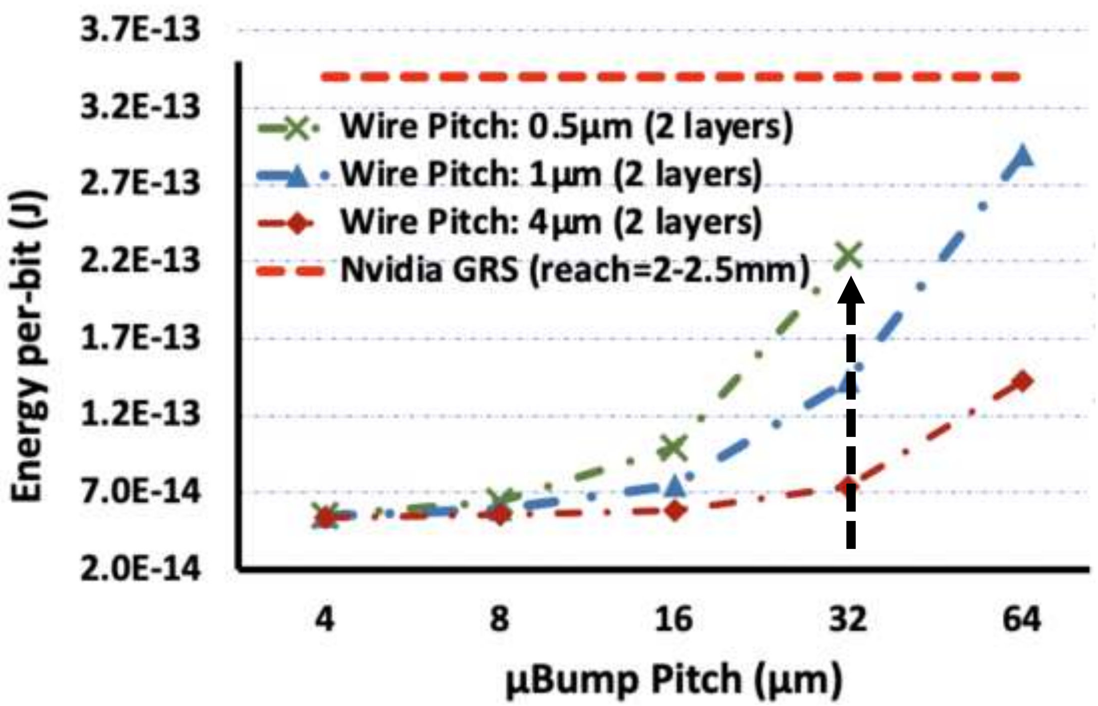

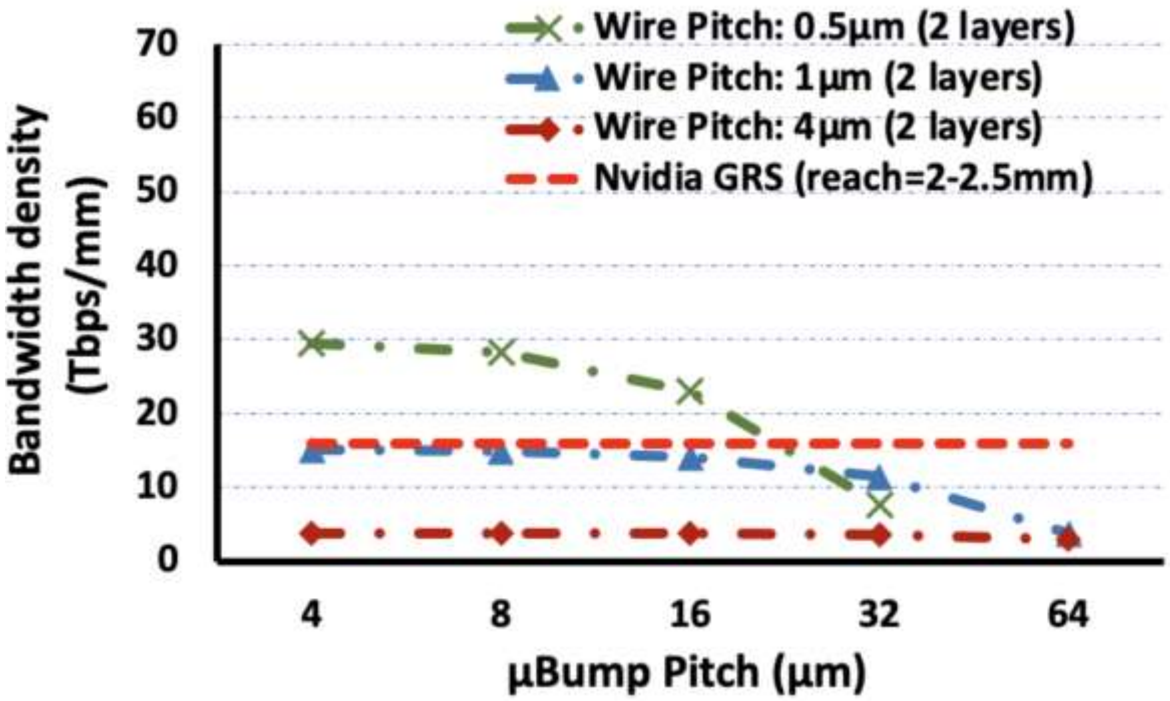

*(1) **Scaling the wire pitch should be accompanied with µbump pitch scaling***

*(2) **Beyond a certain point, benefit of scaling saturates** because of ESD load and minimum wire length*

Impact of ESD-diode Capacitance

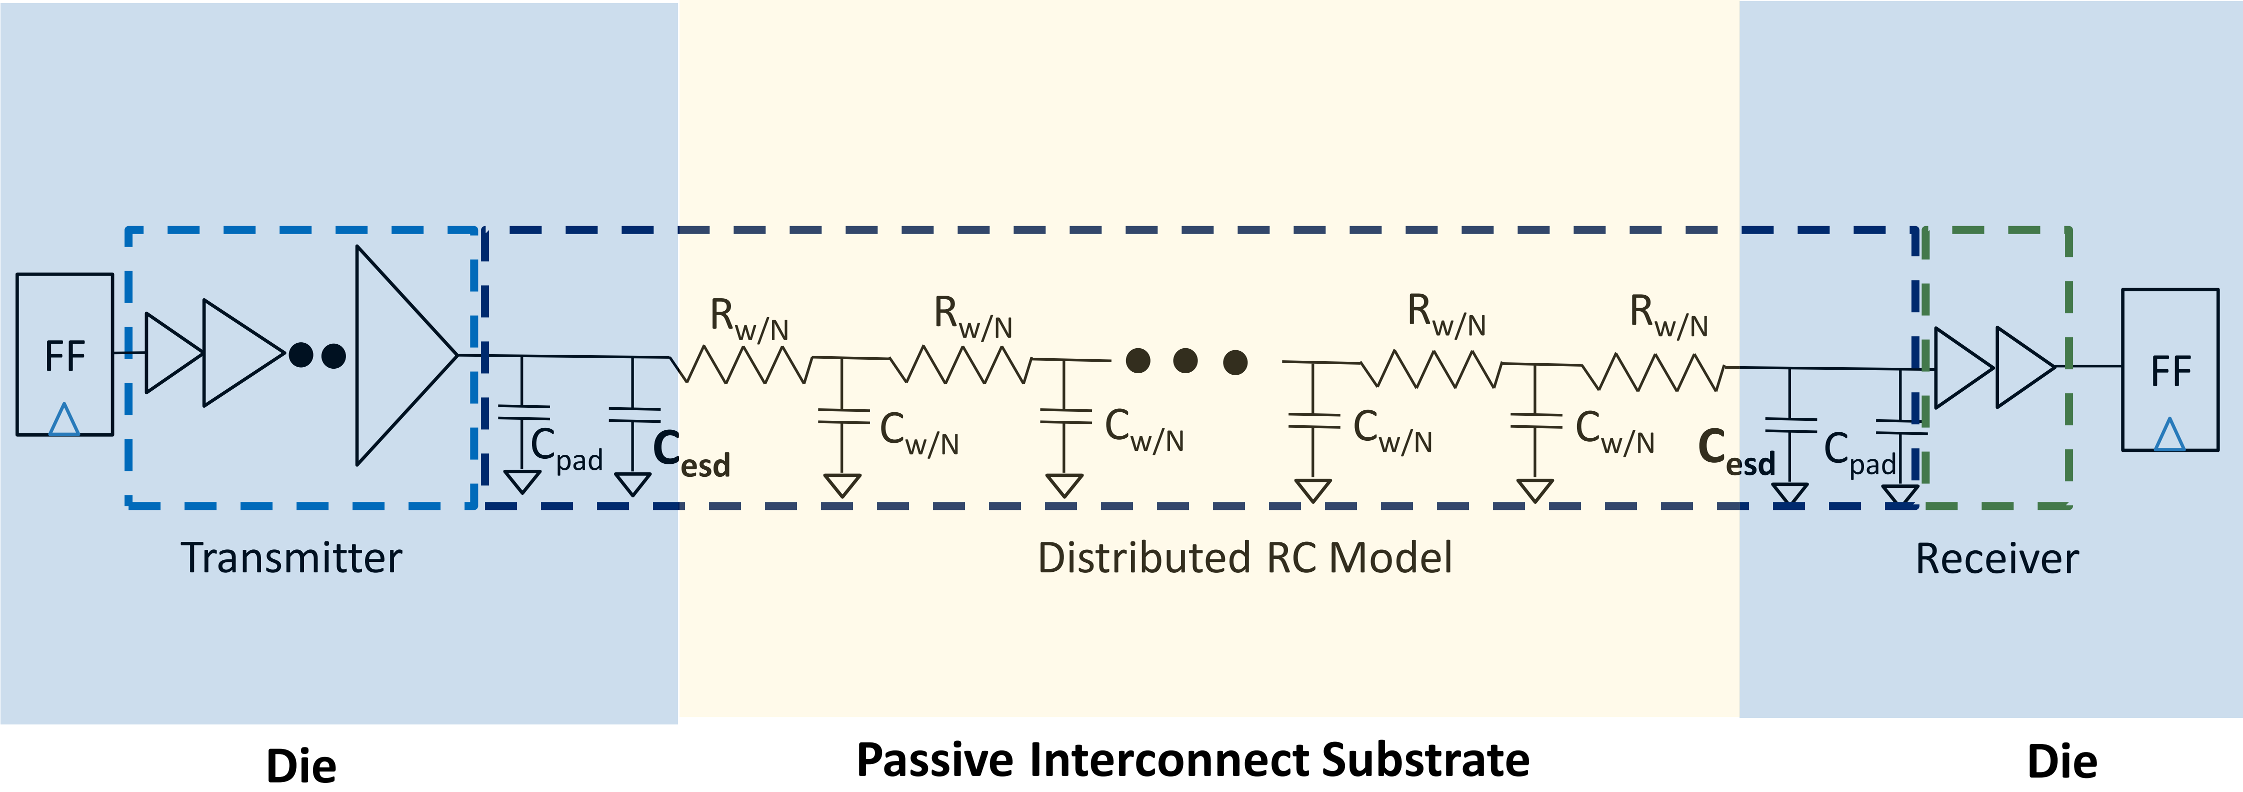

ESD capacitance adds to overall interconnect parasitics and can in fact dominate it

Impact of ESD Capacitance

1 um wire width

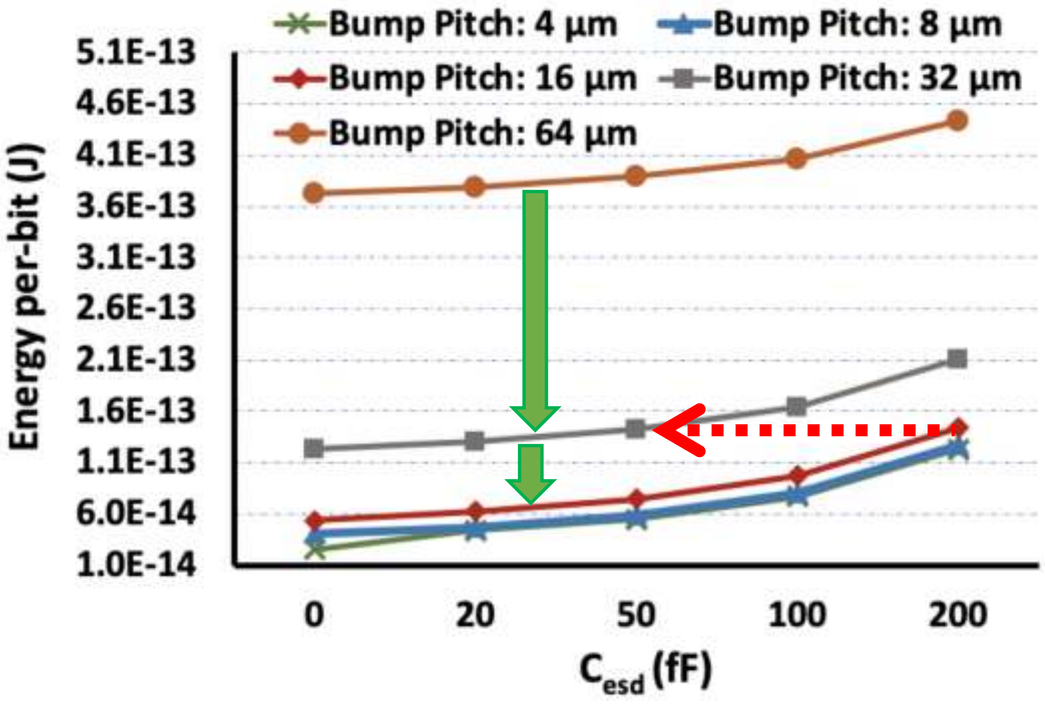

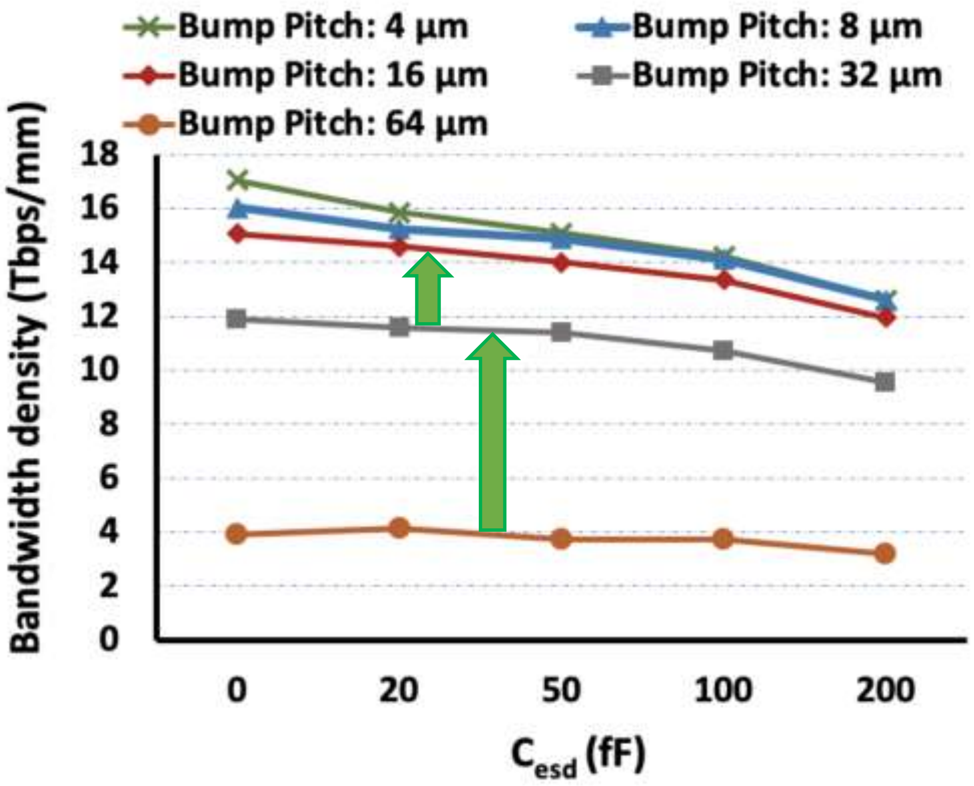

- Reducing ESD-diode capacitance can have the same effect as of reducing ***µBump*** pitch
- ESD-diode capacitance can be used as a lever to scale both energy per bit and bandwidth

Impact of Inter-die Spacing and Dicing Overhead

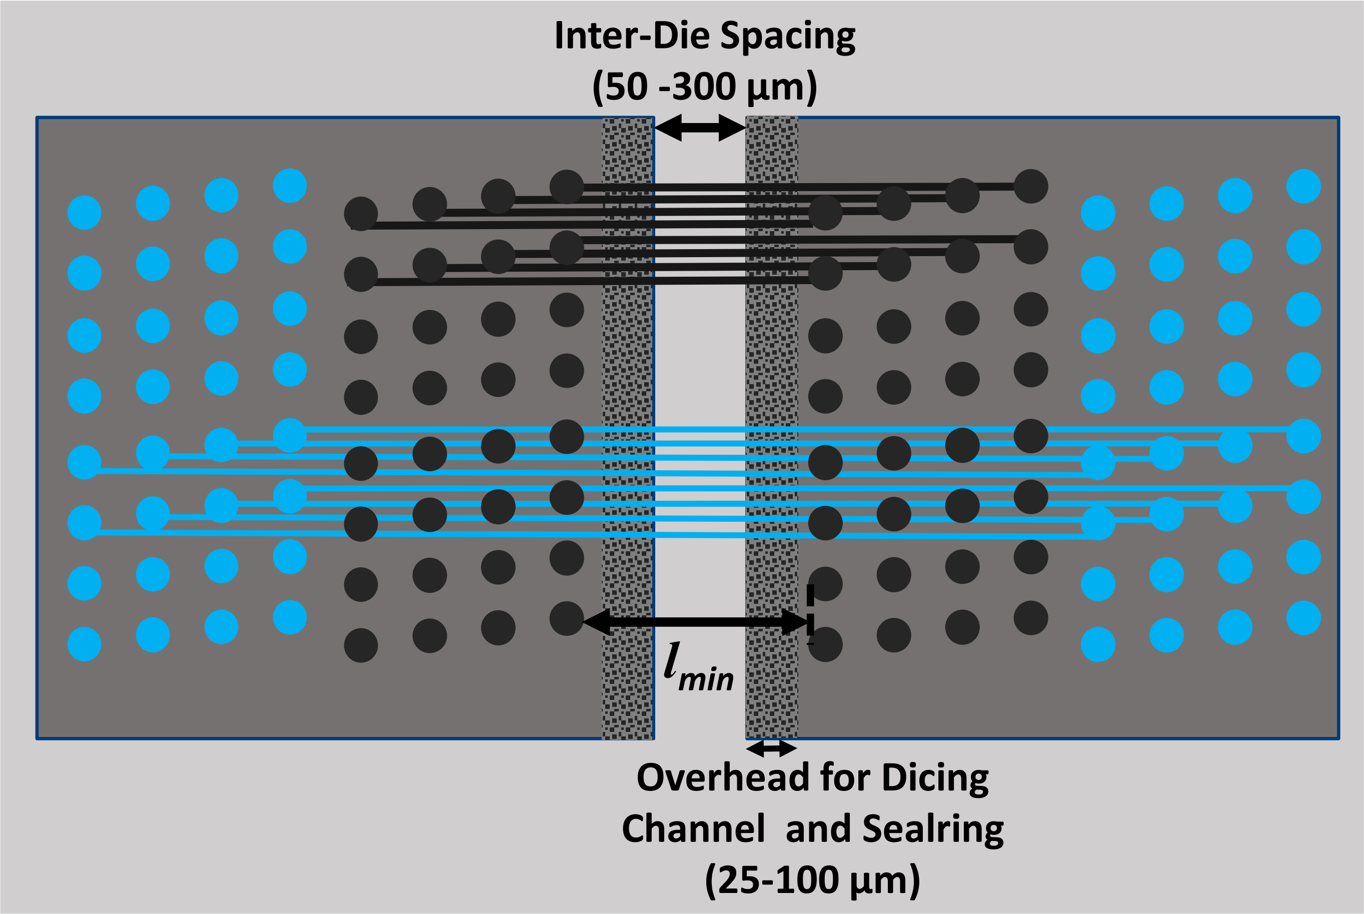

Advanced dicing and die placement technologies can reduce l_{min}

Impact of Inter-die Spacing and Dicing Overhead

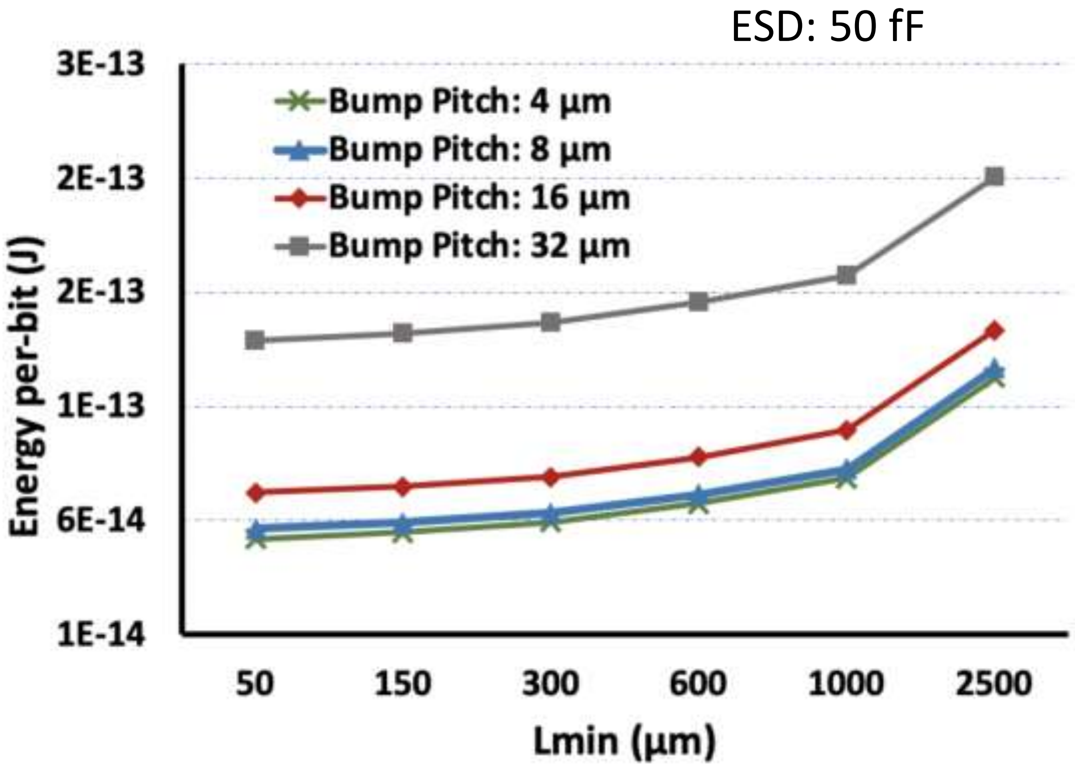

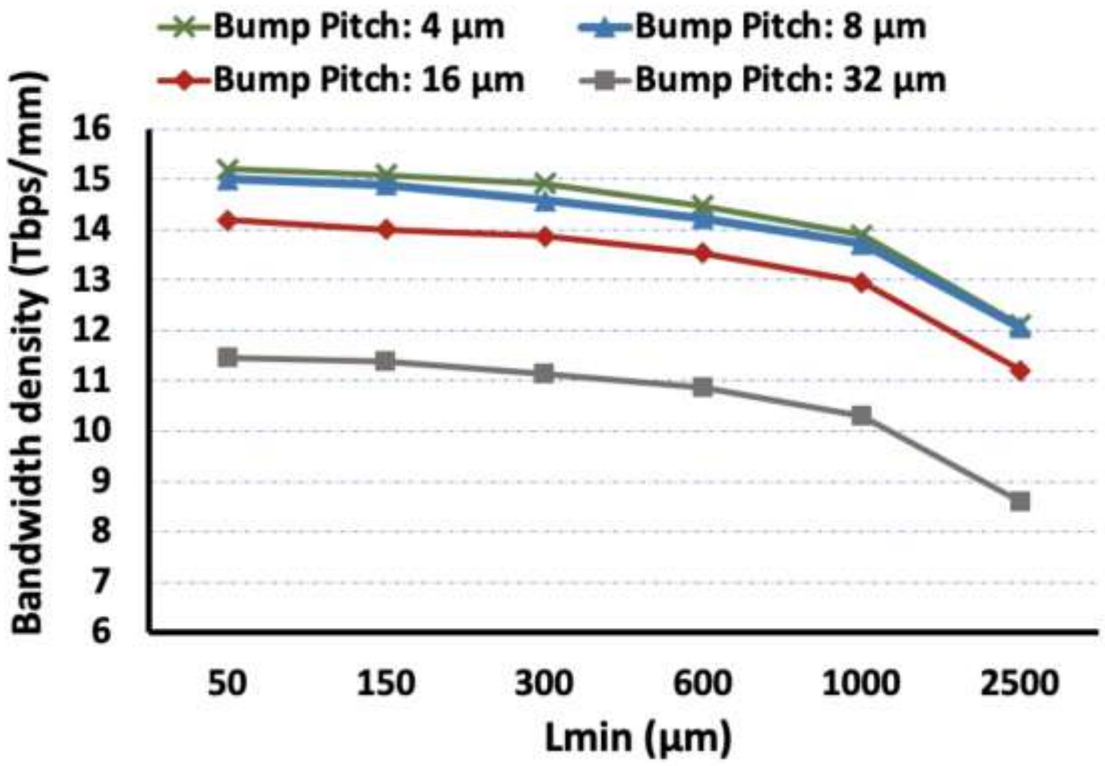

- **Reducing l_{min} below 300 µm provides small improvements**
- **Current generation dicing and placement technologies may be good enough**

Technology Selection in a Heterogeneous Chiplet Eco-system

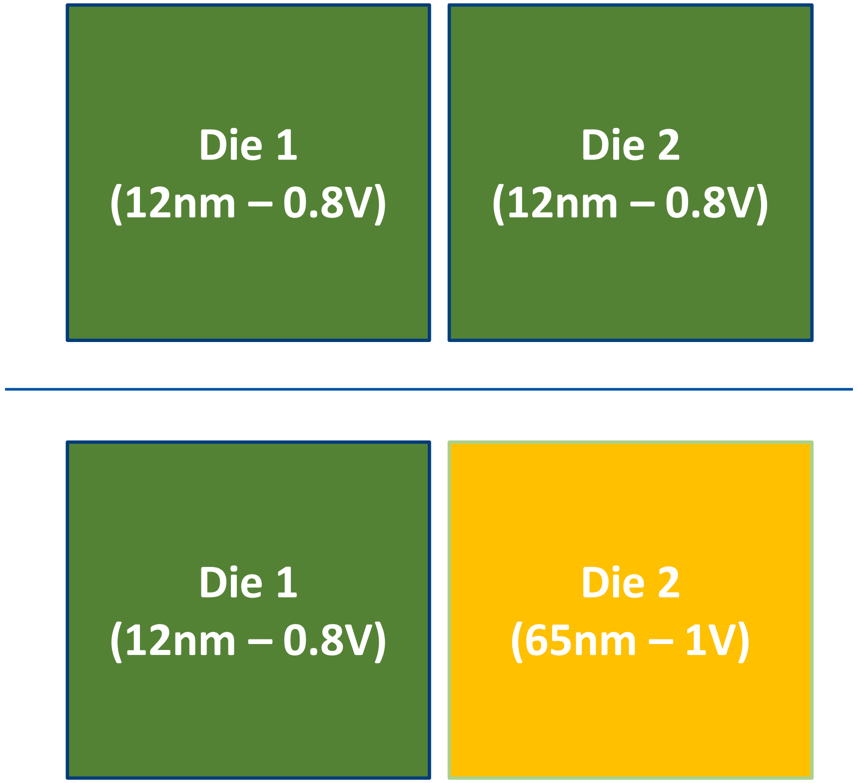

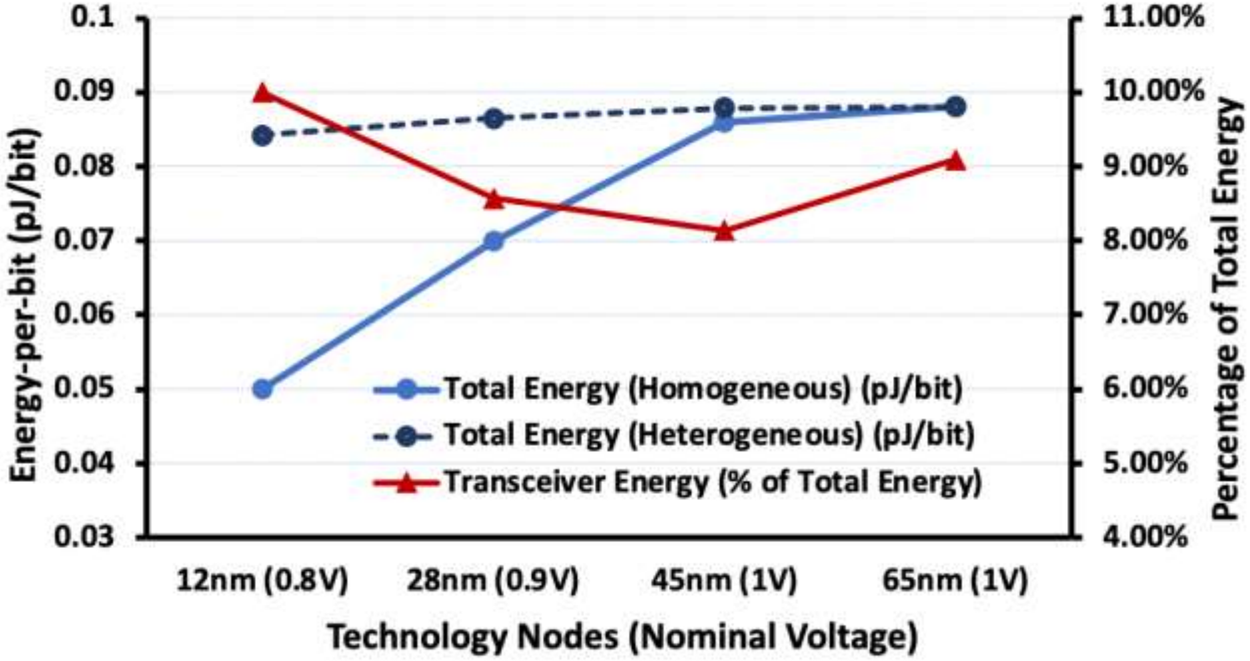

Link efficiency requirements may need to limit the technologies supported by a chiplet ecosystem
Mismatched voltage levels can also have reliability implementations

Conclusions

Chiplets offer THREE primary advantages at the system-level

1. *Heterogeneity.*
   - Integrate logic non-compatible memories + network interfaces
   - Selective upgrades to system IPs
2. *Scale*
   - Ability to build large systems without yield concerns
   - Possibility to go beyond reticle size boundaries (e.g., waferscale)
3. *(Lower-cost) Customization*
   - May be we will see chiplet system variety comparable to board-level systems in near future…

But

- Cost benefits are suspect for high volume, moderately sized SoCs
- Ecosystem challenges remain
- “Wildly” heterogeneous chiplet systems may not be a good idea

Backup

Thank you sponsors!

ADVANTEST®

Amkor Technology®

synopsys®

MEPTEC
THE NEXT GENERATION

Global Companies Rate Advantest THE BEST ATE Company 2021

Advantest receives highest ratings from customers in annual VLSIresearch Customer Satisfaction Survey for 2 consecutive years.

Global customers name Advantest THE BEST supplier of test equipment in 2020 and 2021, with highest ratings in categories of:

Technical Leadership – Partnership – Trust – Recommended Supplier – Field Service

"Year-after-year the company has delivered on its promise of technological excellence and it remains clear that Advantest keeps their customers' successes central to their strategy. Congratulations on celebrating 33 years of recognition for outstanding customer satisfaction."

— **Risto Puhakka**, President VLSIresearch

Amkor's Differentiators

Technology

Advanced Packaging Leadership
Engineering Services
Broad Portfolio

Quality

QualityFIRST Culture
Execution
Automation

Service

Design & Test Through Drop Ship
Manufacturing Footprint
Local Sales & Support

SYNOPSYS®

Silicon to Software™

COPYRIGHT NOTICE