



THE CHIPS TO SYSTEMS CONFERENCE

SHAPING THE NEXT GENERATION OF ELECTRONICS

JUNE 23-27, 2024

MOSCONE WEST CENTER
SAN FRANCISCO, CA, USA





SHAPING THE NEXT GENERATION OF ELECTRONICS

JUNE 23-27, 2024

MOSCONE WEST CENTER
SAN FRANCISCO, CA, USA

SYNOPSYS®

Modeling Networks-on-Chip for Architecture Analysis and Optimization


Kamal Desai

Tim Kogel

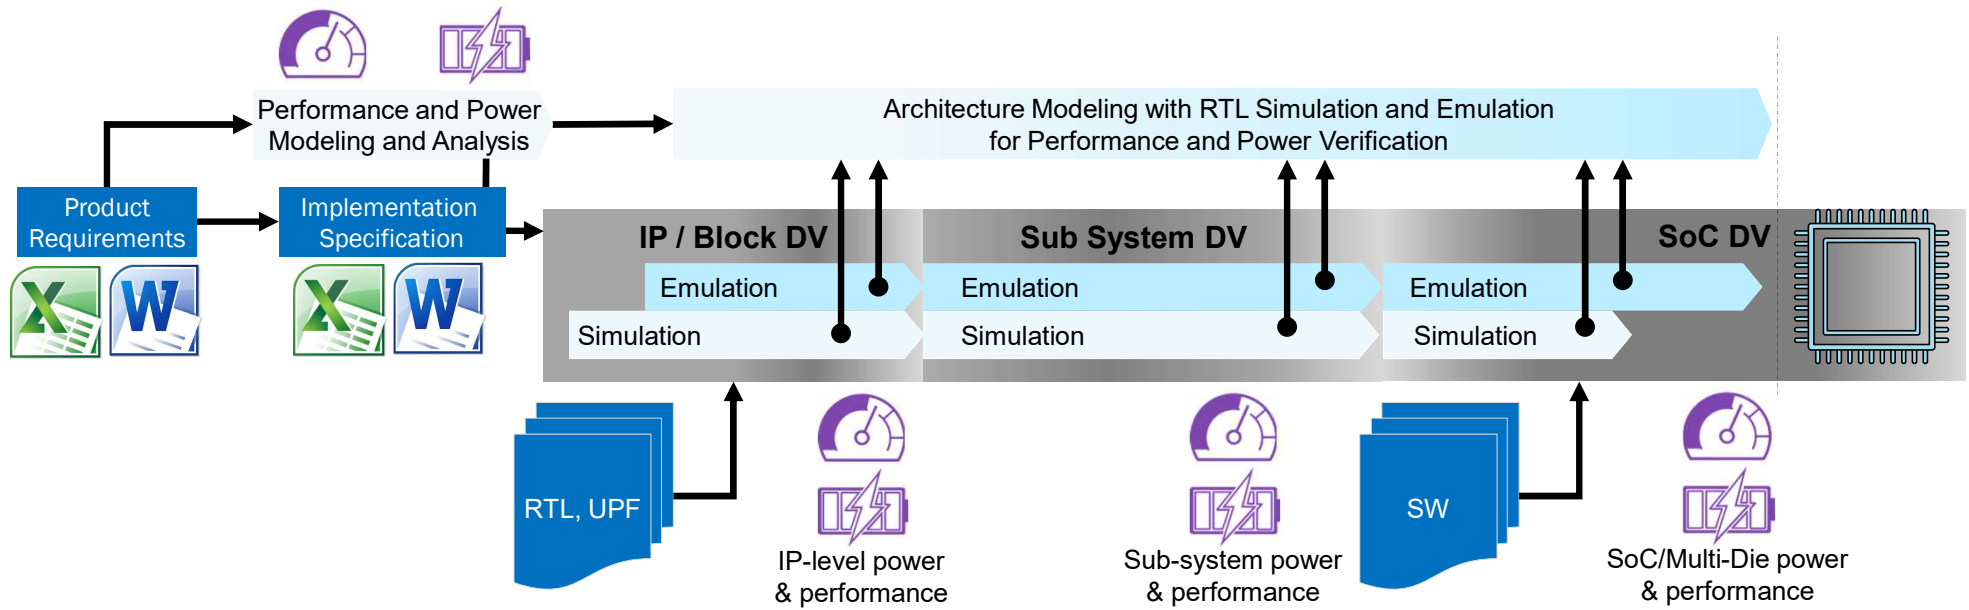


“If you think good architecture is expensive, try bad architecture”

Brian Foote and Joseph Yoder, Department of Computer Science, University of Illinois at Urbana-Champaign, <http://www.laputan.org/mud>, 1999



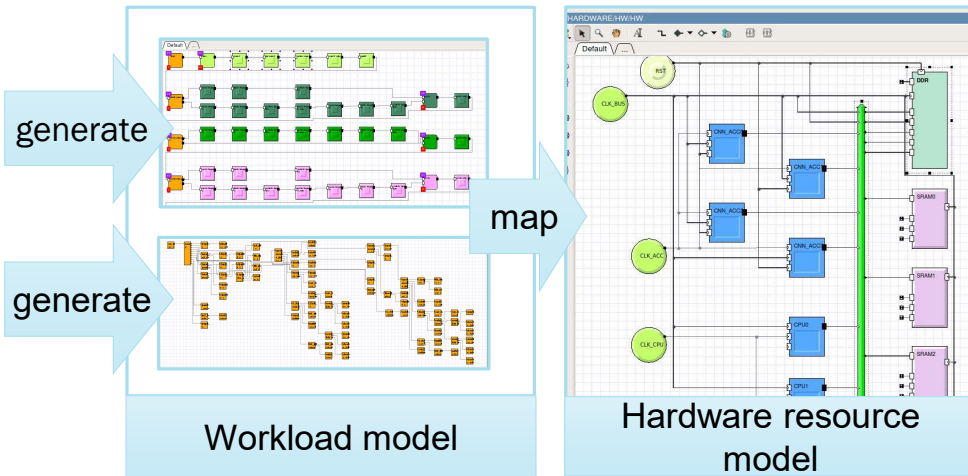
Designing the Right Architecture



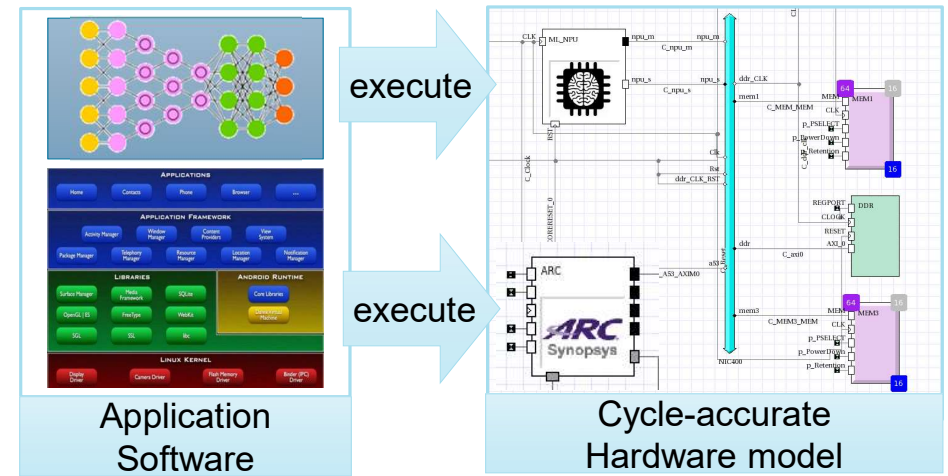
- **Goal:** design the right product, de-risk architecture, translate product requirements into implementation specification, track requirements
- **Use cases:** KPI-driven power/performance analysis, HW/SW partitioning, interconnect/memory optimization, cache/cache-coherency dimensioning

Architecture Analysis with Virtual Prototyping

Early architecture exploration and optimization with workload models



Performance validation with Software



Architecture Analysis with Virtual Prototyping


Early architecture exploration and optimization with workload models

- KPI capture and sensitivity analysis
- Traffic and application workload modeling
- HW/SW partitioning, architecture specification
- Multi-chip power/performance analysis

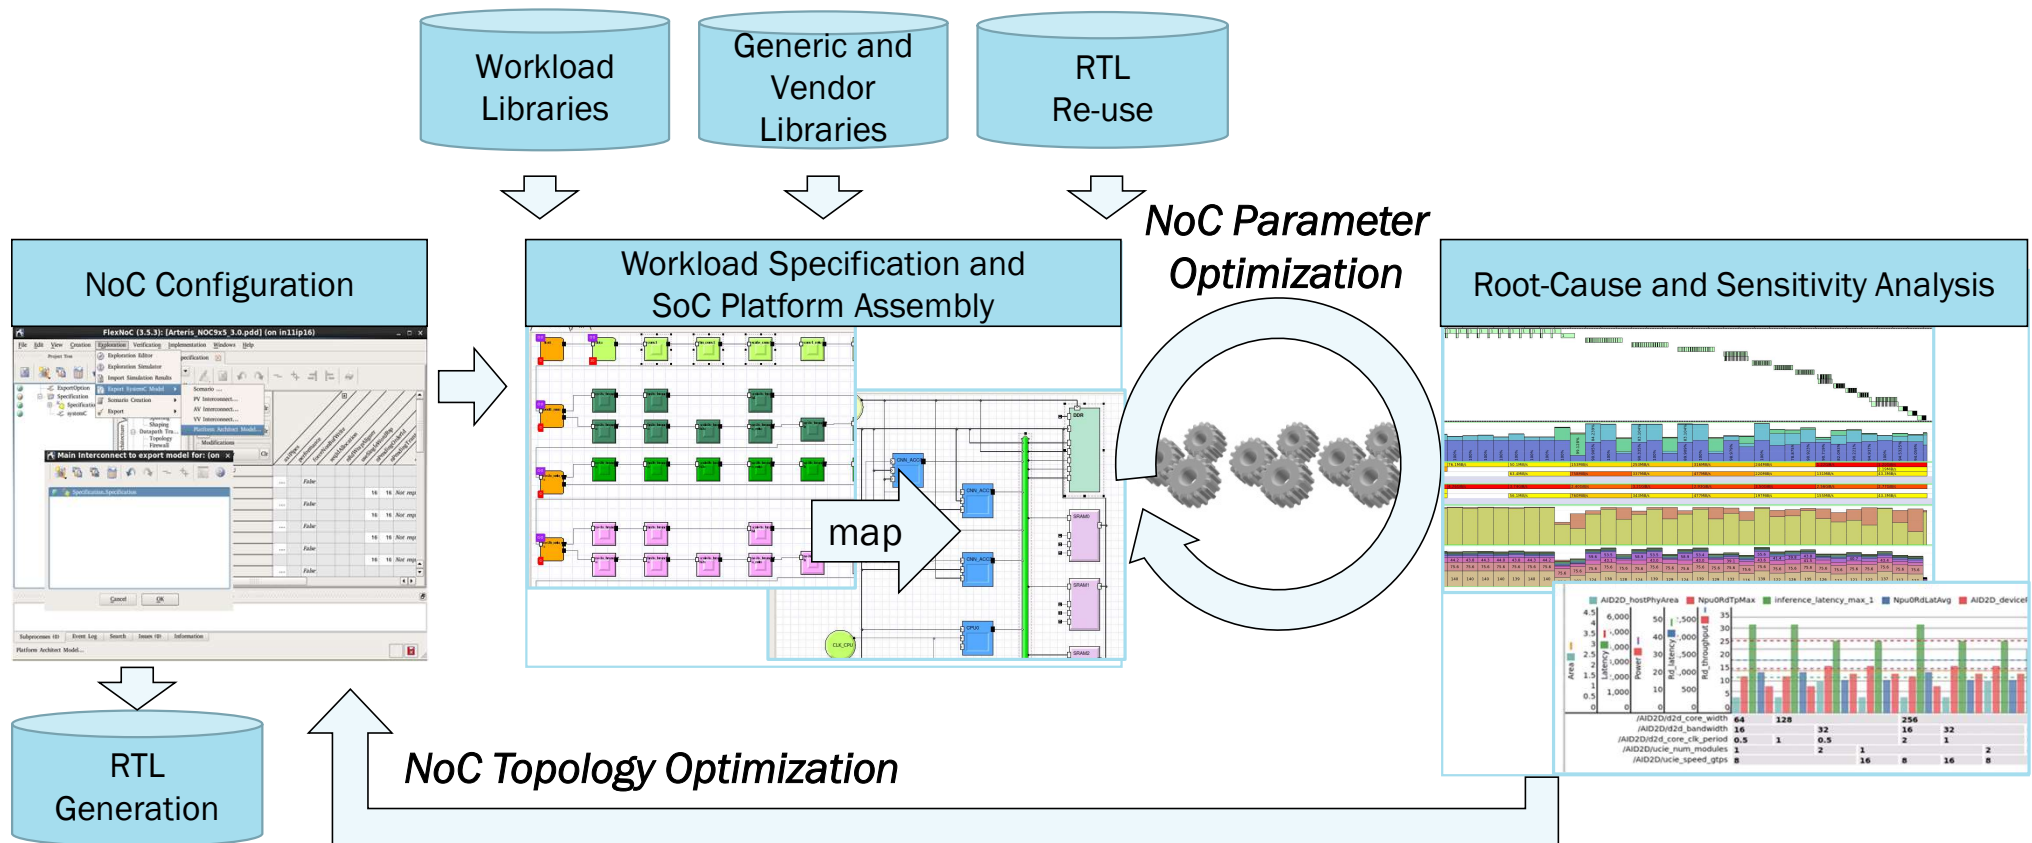
Performance validation with Software

- KPI tracking and validation
- IP optimization and benchmarking
- SoC performance validation
- L1/L2 cache & cache coherency optimization

- Interconnect/memory performance optimization
 - L3 cache & cache coherency optimization
- System-level power/energy analysis and optimization



Networks-on-Chip Analysis and Optimization



Trend: Multi-Die Ecosystem

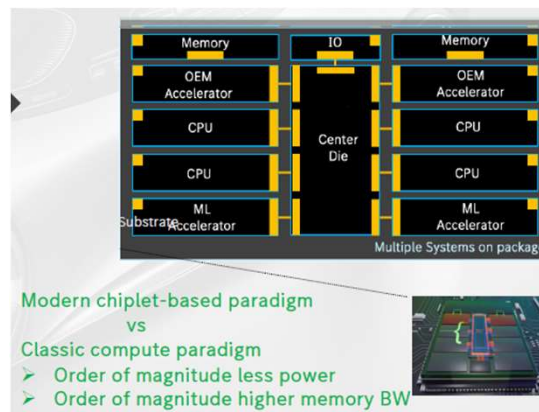
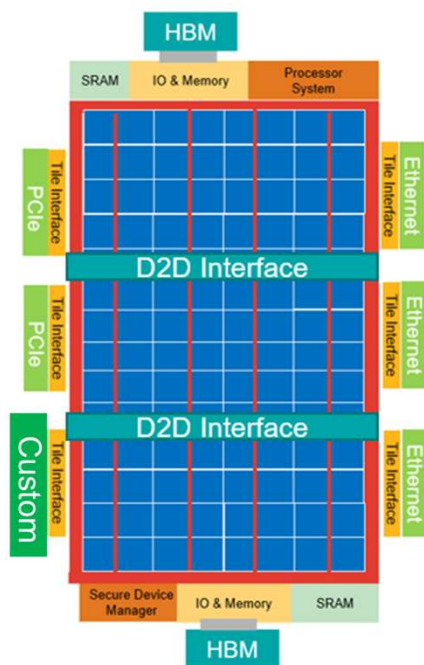
Specialization with domain-specific accelerator chiplets

"Data Driven Design for Adaptive Multi-Die Systems"

Vikrant Kapila, Principal Engineer & Director, Altera
Synopsys Users Group Conference, Silicon Valley, 2024

"Standardization for Automotive Computing Hardware"

François Piednoël, Distinguished Architect, Mercedes-Benz
Automotive Compute Conference, US, 2024



"Standardization could open door to third-party chiplets"

AMD's CTO Mark Papermaster and SVP Sam Naffziger

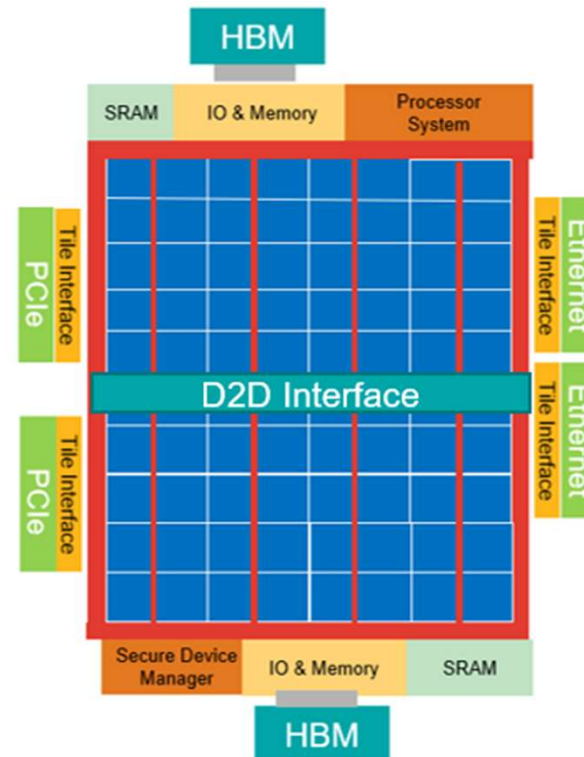
https://www.theregister.com/2024/03/27/amd_chiplets_future

- Baseline compute, memory, and IO dies with standard multi-die interface
- Custom companion chiplets with access to memory infrastructure
- Architecture challenges:
 - Workload partitioning & mapping
 - Die-to-die interconnect dimensioning
 - End-to-end performance, power, and thermal analysis
 - Physical design planning

Intel® Agilex FPGA: Data Driven DSE

Synopsys Platform Architect: System-C Interconnect + Peripherals:

- HNOG with fabric facing INIUs
 - SystemC TLM 2.0 Interconnect Model offers same set of parameters and Design Space Exploration as hardware.
 - VCD based topology debug for congestion.
- uNOC direct to m20ks
 - System TLM 2.0 inhouse model.
- 3rd Party DDR Controllers
 - SystemC TLM 2.0 3rd party model.
- HBM Controllers
- Inhouse correlation done to RTL.
- Hard Processor Sub-system
 - Cache Coherent Interconnect
 - Real Time processor & Application processor subsystem
- Core Fabric
 - DSP sub-system, AI tensor Block DSP.
- Memory Access Patterns (TGs)
 - Use-case workloads as DFGs
 - 1 D traffic generators, trace file-based traffic



Intel Agilex® Chiplet Model ecosystem

System Workload Mapping

Early Workload Design Space Exploration using Synopsys Platform Architect

Embedded: e.g., FFT





Fig. 2. Binary tree representation. (a) Node split in a binary tree. (b) Binary tree for a 32-point radix-2 DIF FFT. (c) Binary tree for a 32-point radix-2³, 2² DIF FFT.

$$F(x, y) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n) e^{-j2\pi(x\frac{m}{M} + y\frac{n}{N})}$$

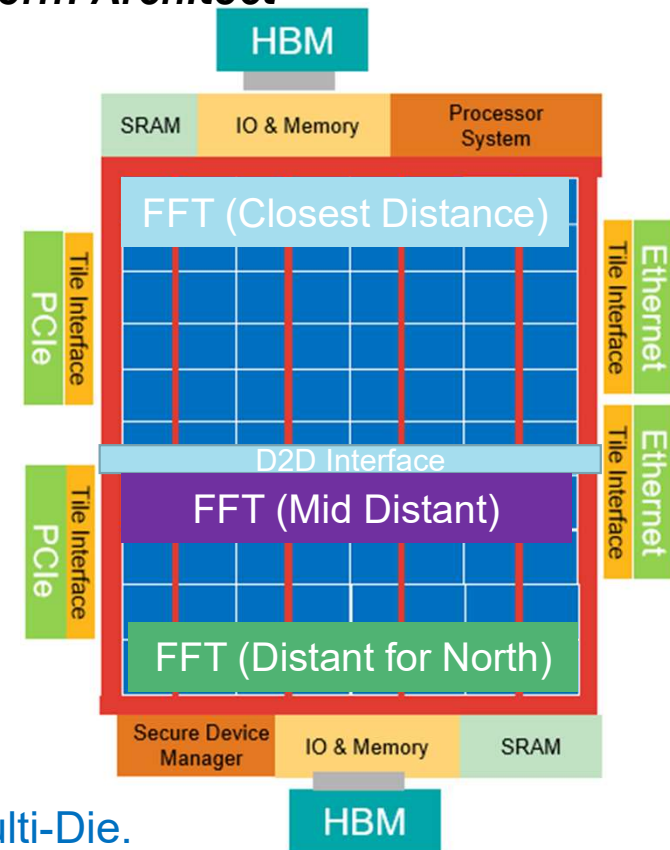
$$f(m, n) = \frac{1}{MN} \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} F(x, y) e^{j2\pi(x\frac{m}{M} + y\frac{n}{N})}$$

MAP

Eff. Utilization comparison chart

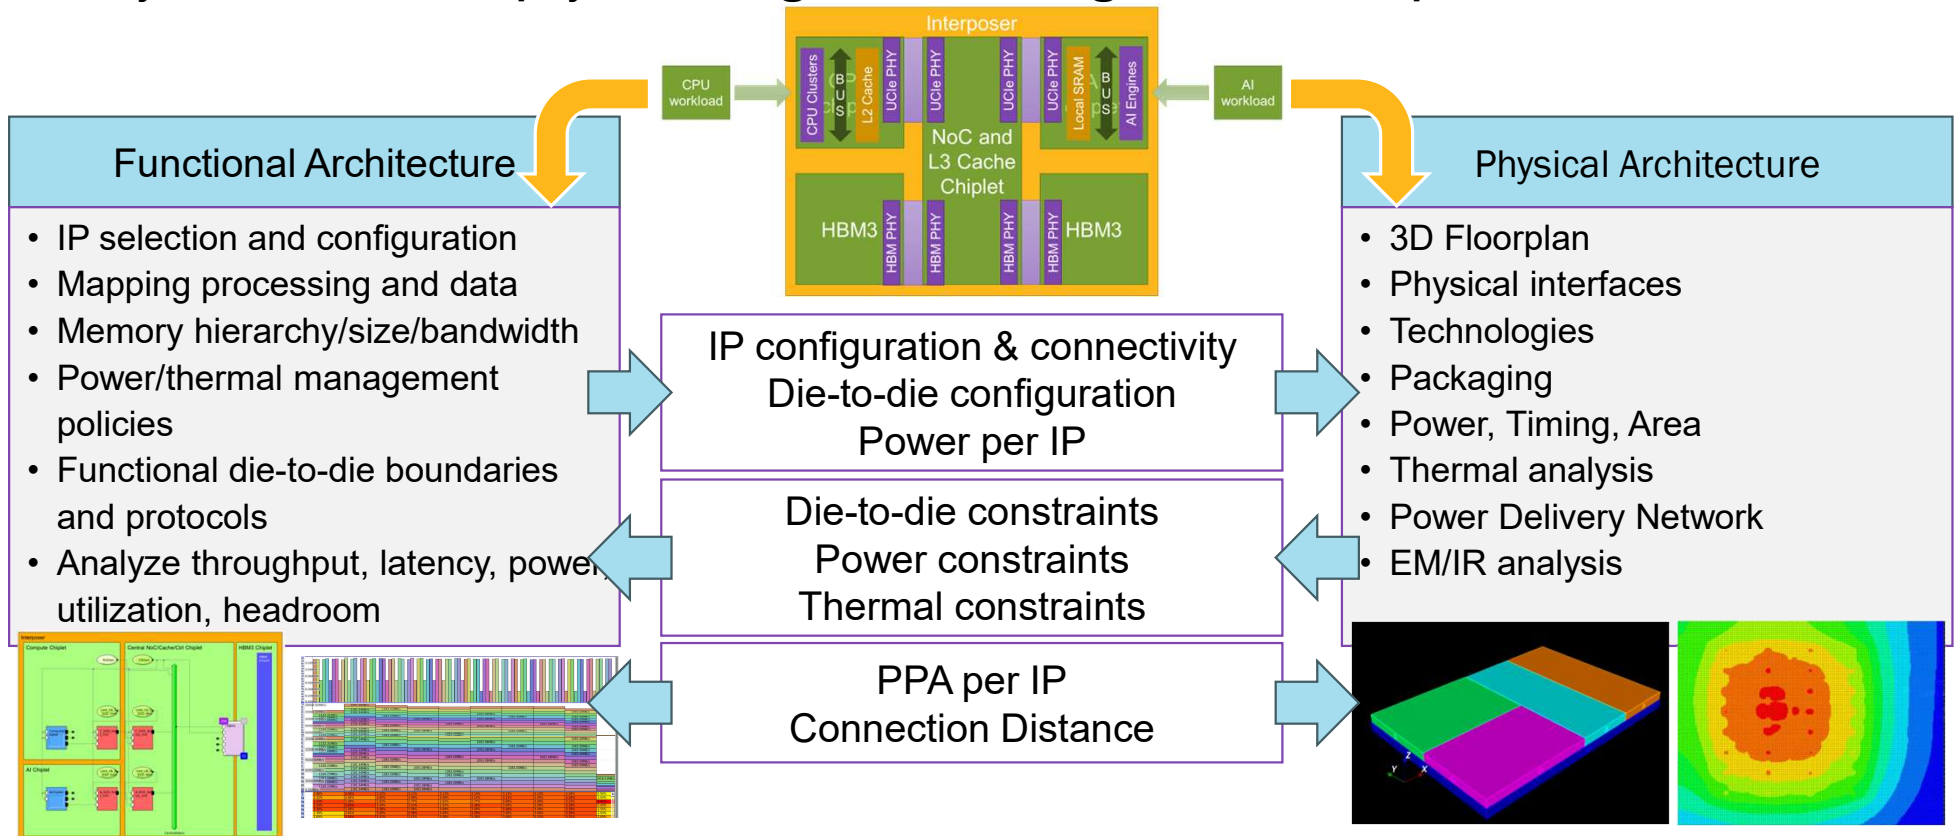


Monolithic single u Processor approach for FFT's doesn't scale for Multi-Die.




Trend: Physical-aware Architecture Optimization

Early consideration of physical design issues during architecture specification



Summary

Data-driven Networks-on-Chip Analysis and Optimization



- Early exploration and analysis visibility with orders of magnitude better turnaround times



- Enable QoS analysis and optimization for your application traffic and topology



- Achieve system performance and power goals without costly over or under designs

THANK YOU !