SYNOPSYS®

Opportunities for Optics into the Datacenter and High-Performance Computing (HPC) Infrastructure

Accelerating PIC Adoption in Established Markets

Jigesh Patel, Remco Stoffer*

April 17, 2024 – PIC International – Brussels (Belgium)

CONFIDENTIAL INFORMATION

The information contained in this presentation is the confidential and proprietary information of Synopsys. You are not permitted to disseminate or use any of the information provided to you in this presentation outside of Synopsys without prior written authorization.

IMPORTANT NOTICE

In the event information in this presentation reflects Synopsys' future plans, such plans are as of the date of this presentation and are subject to change. Synopsys is not obligated to update this presentation or develop the products with the features and functionality discussed in this presentation. Additionally, Synopsys' services and products may only be offered and purchased pursuant to an authorized quote and purchase order or a mutually agreed upon written contract with Synopsys.

© 2024 Synopsys, Inc.

Overview

- Introduction
  - Motivation
- Optics in Data Centers
  - Emergence of AI/ML, Explosion of Data Traffic
  - Flattening of Networks for Low Latency, and Aggregation of Homogeneous Resources
  - NPO, CPO
- Re-thinking Electrical-Optical Interface
  - Optical I/O Roadmap
- Optics as an Enabler to PCIe and UCIe
  - “Convergence” of Ethernet, PCIe and UCIe
- 3D-HI and Impact on the Design Automation
  - Multi-Die, System-of-Chips Implications
- Summary

© 2024 Synopsys, Inc.

Acronyms in Order of Appearance

HPC: High-Performance Computing
PIC: Photonic Integrated Circuit
AI: Artificial Intelligence
ML: Machine Learning
I/O: Input and Output
NPO: Near-Packaged Optics
CPO: Co-Packaged Optics
PCIe: Peripheral Component Interconnect Express
UCIe: Universal Chiplet Interconnect Express
3D-HI: 3D Hybrid Integration
PC: Personal Computer
AR/VR: Augmented/Virtual Reality
5G: 5th Generation
LiDAR: Light Detection & Ranging
E-O-E: Electrical, Optical, Electrical
NRZ: Non-Return to Zero
PAM4: 4-level Pulse Amplitude Modulation
QPSK: Quadrature Phase-Shift Keying
HBM: High-Bandwidth Memory
MCM: Multi-Chip Module
TOR, ToR: Top-of-the Rack
DAC: Digital-to-Analog Converter
GPU: Graphic Processing Unit
ISP: Internet Service Providers
DC: Data Center
DCI: Data Center Interconnect
POD: Not an acronym (modular and scalable unit)
CPU: Central Processing Unit
QSFP: Quad Small Form-Factor Pluggable
BGA: Ball Grid Array
OE: Optoelectronic
SERDES: Serializer, De-Serializer
LR: Long Reach
VSR: Very Short Reach
PHY: Physical layer (of the networking model)
DSP: Digital Signal Processing
DRV: Driver
TIA: Trans-Impedance Amplifier
PD: Photo-detector
Mod: Modulator, Modulation

© 2024 Synopsys, Inc.

Acronyms in Order of Appearance (Continued)

IC: Integrated Circuit
CXL: Compute Express Link
Tx, TX: Transmitter
Rx, RX: Receiver
tp#: Test Point #
IP: Intellectual Property
IBIS: Input/Output Buffer Information Specification
AMI: Algorithmic Modeling Interface
SnP: Data file extension for a Touchstone file
PDK, ADK: Process/Assembly Design Kit
XSR: Extra-short Reach
AFE: Analog Front End
MZM: Mach-Zehnder Modulator
MRM: Micro-Ring Modulator
EAM: Electro-Absorption Modulator
BW: Bandwidth
RIN: Relative Intensity Noise
MB/s, GB/s, TB/s: Mega, Giga, Tera bits per second
SIP: Silicon Photonics
xPU: Cross-Point Unit
D2D: Die-to-Die
OIF: Optical Internetworking Forum
PCI: Peripheral Component Interface
SIG: Special Interest Group
GbE, TbE: Giga-bit and Tera-bit Ethernet
ZR: Zero-dispersion shifter
LR: Long Range
EDA: Electronic Design Automation
SW: Software
Pkg: Package
DDR: Double-Data Rate
InFO: Integrated Fan-Out
RDL: Re-Distribution Layer
KGD: Known Good Die
TCAD: Technology Computer Aided Design
RF: Radio Frequency
EM: Electro-magnetic
PEX: Parasitic Extraction
DRC: Design Rule Check
LVS: Layout-versus-Schematic

© 2024 Synopsys, Inc.

Introduction/Motivation

Accelerating PIC Adoption in Established Markets

© 2024 Synopsys, Inc.

Changing Market Dynamics

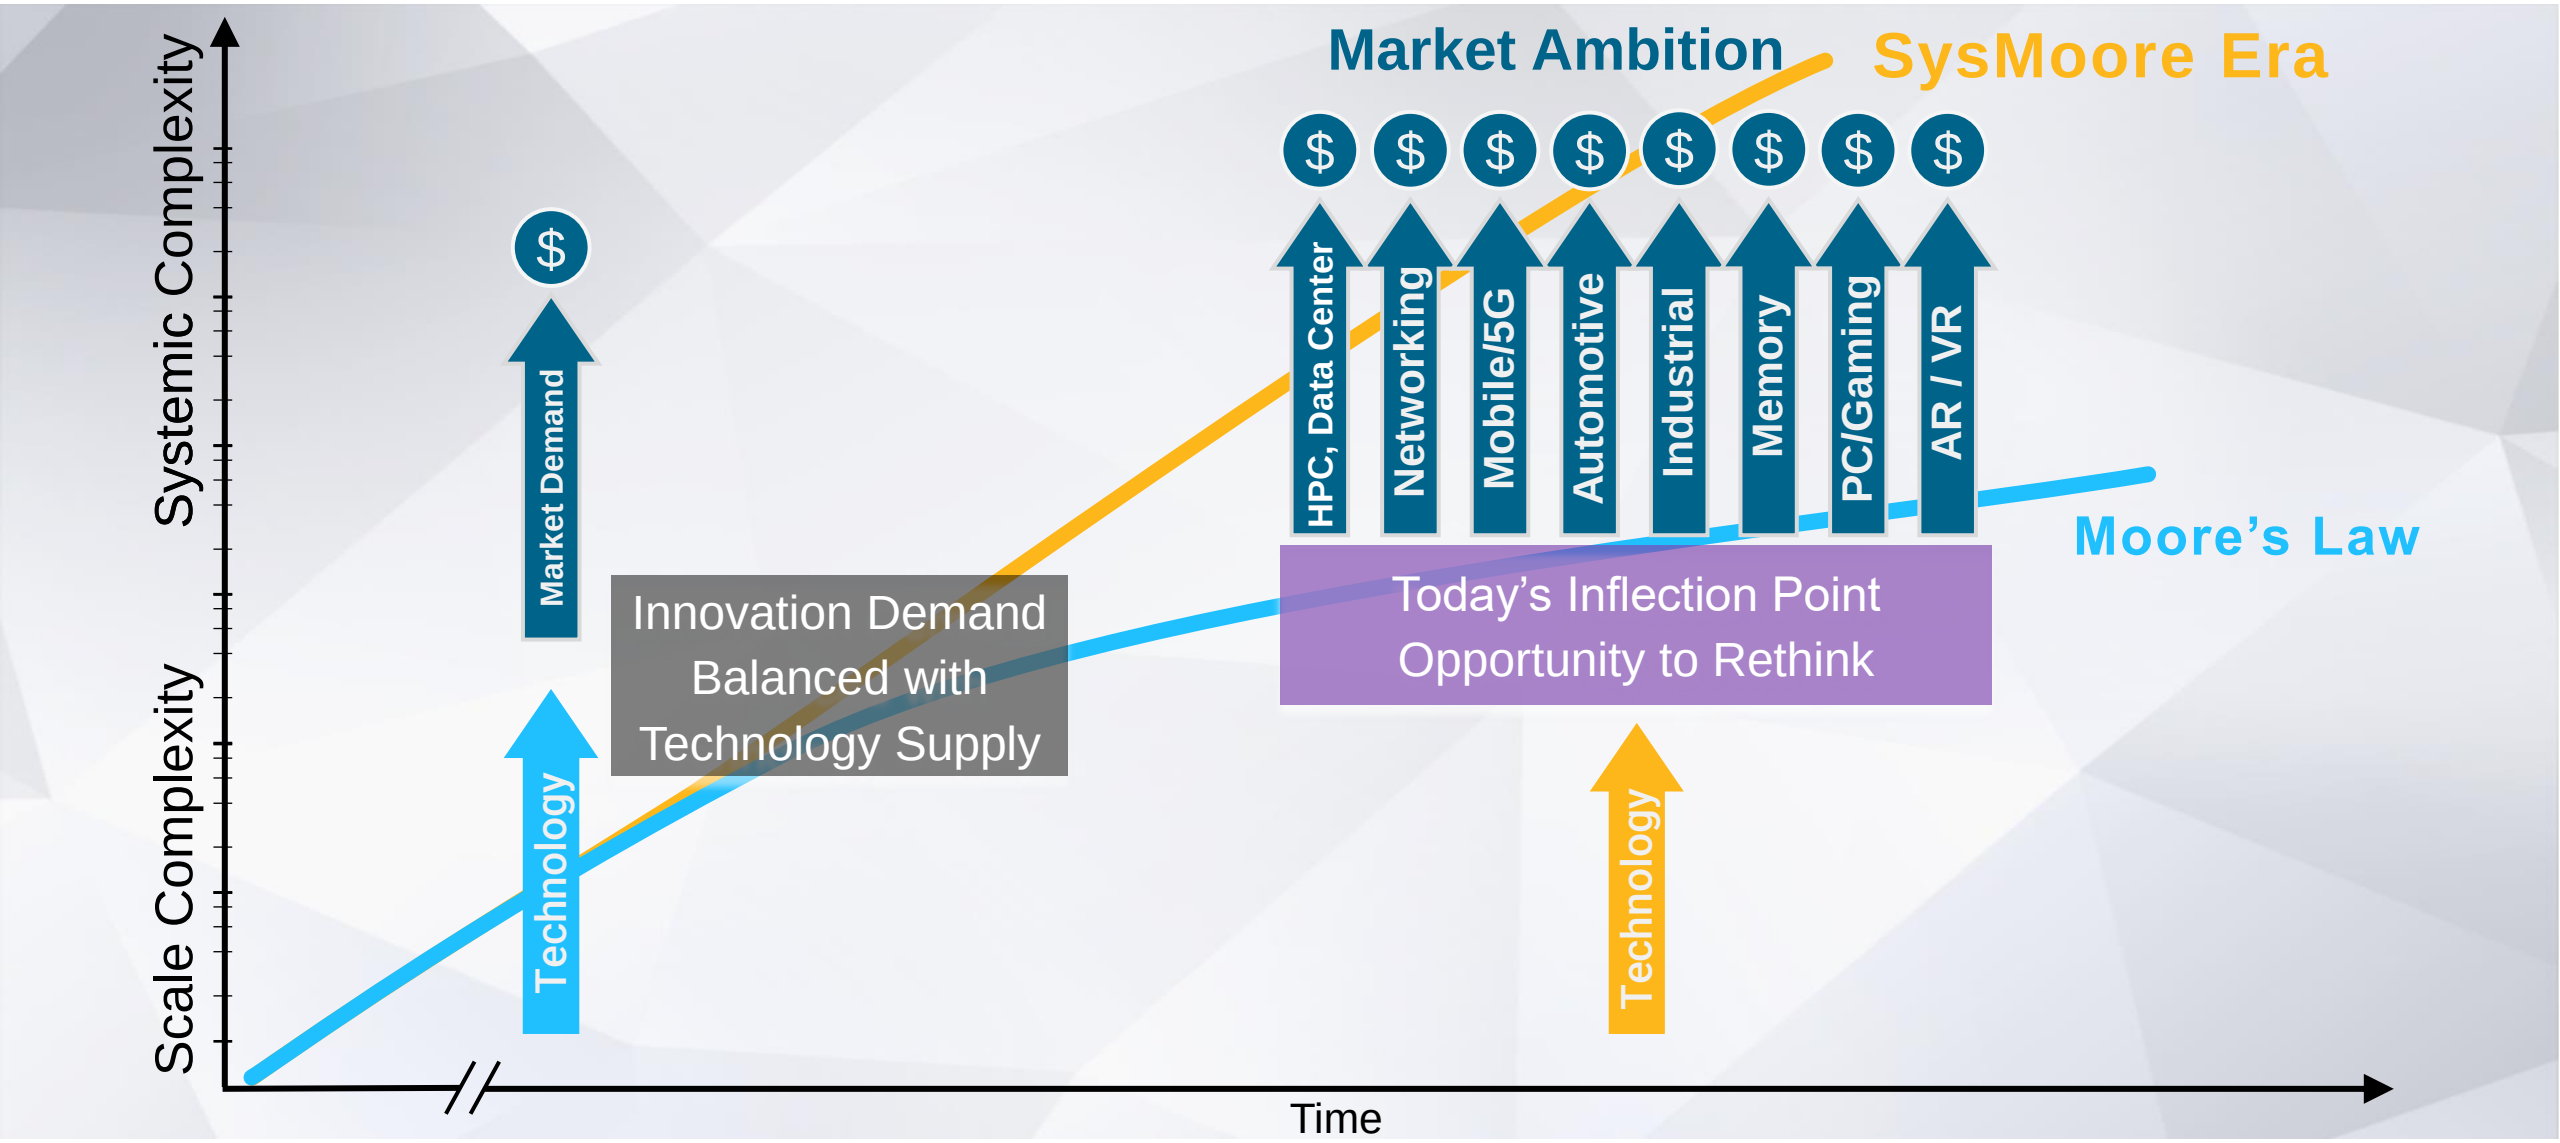

© 2024 Synopsys, Inc.

Technology Updates and Future scope

New applications are emerging because

- Loss handling is getting more expensive in copper and costly with the data rate growing rapidly
- AI/ML demands are growing and requesting different architectures to save power, reduce latency
- Different standards are adopting E-O-E interconnect driven by the need for low power and low latency → PCIe over optics, Optical Chiplets, and more

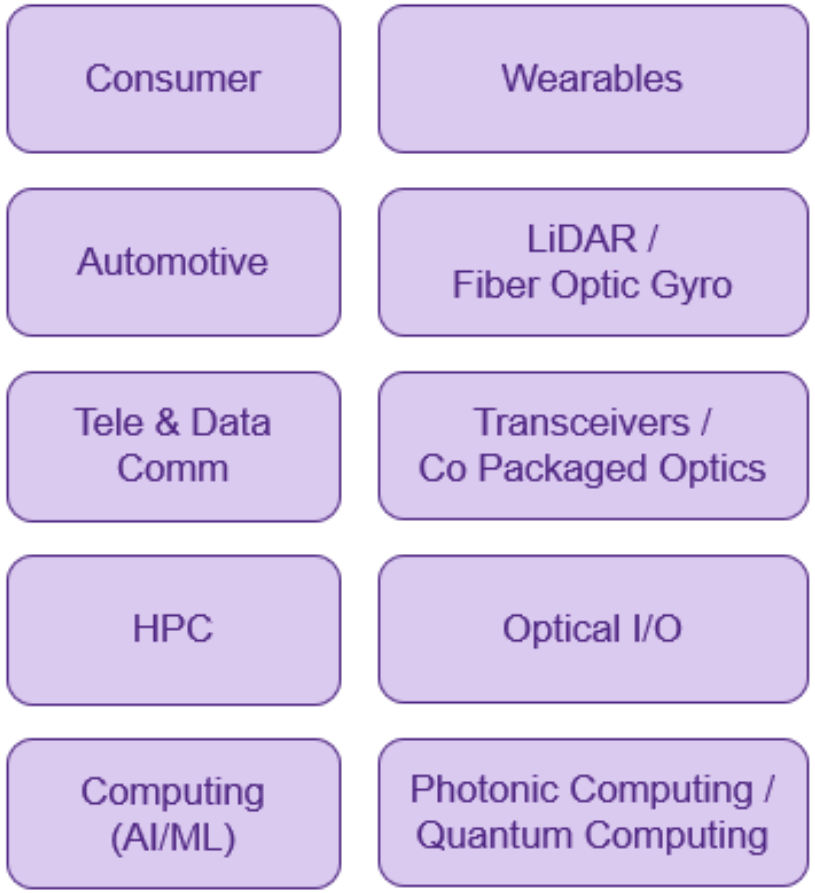

© 2024 Synopsys, Inc.

Opportunities for Optics to go Deeper into the Data Centers

Accelerating PIC Adoption in Established Markets

© 2024 Synopsys, Inc.

Interconnect Scaling Challenges

AI/ML growth has led to an increased cluster size and IO demand

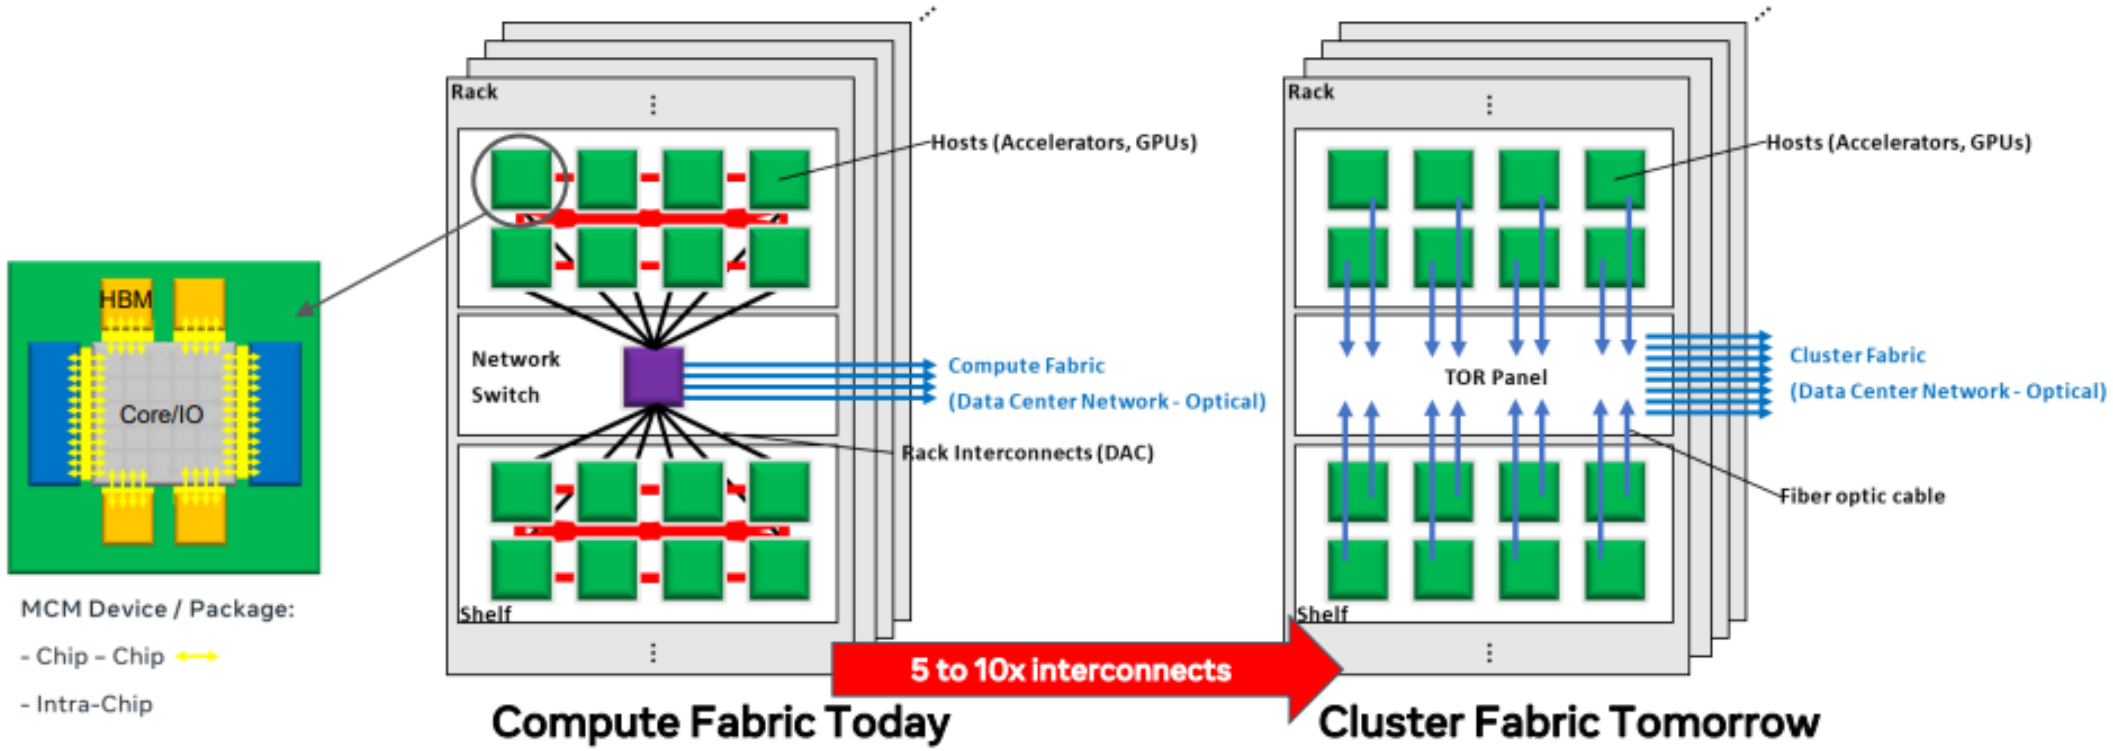

Source: Meta

© 2024 Synopsys, Inc.

Main Trends Driving Optics into the Data Center

"Fatter" links extending up to 10s of meters within pod efficiently supported with optical interconnects

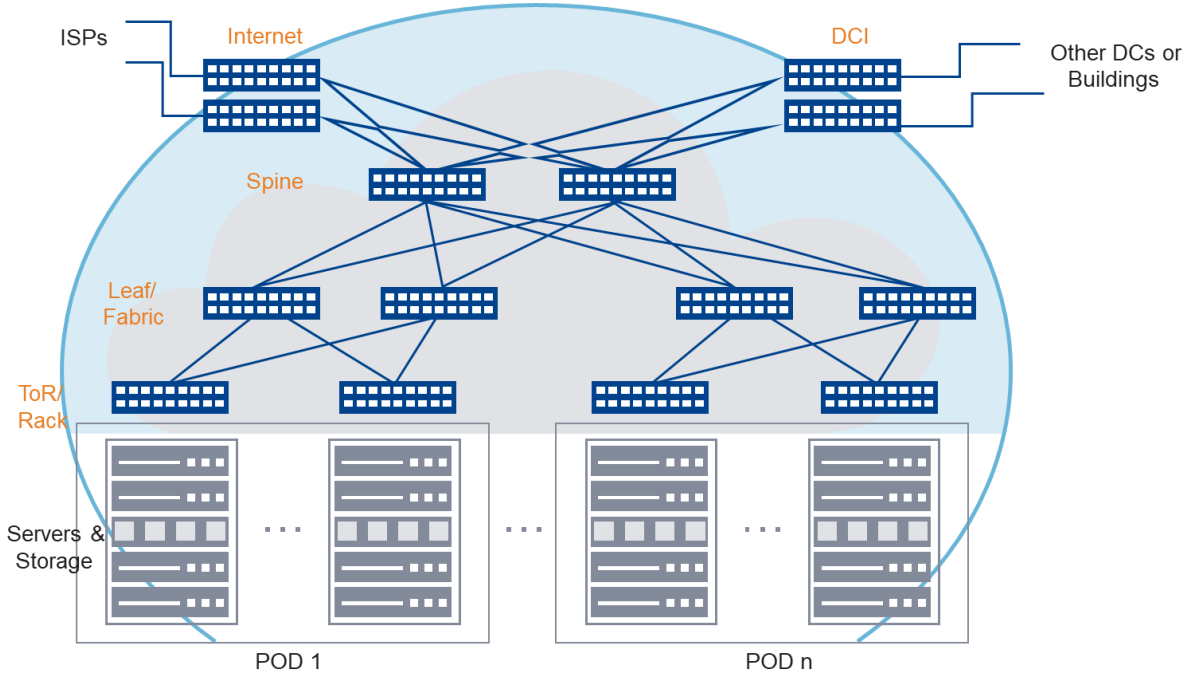

Explosion of Data Traffic in Data Centers

© 2024 Synopsys, Inc.

Main Trends Driving Optics into the Data Center

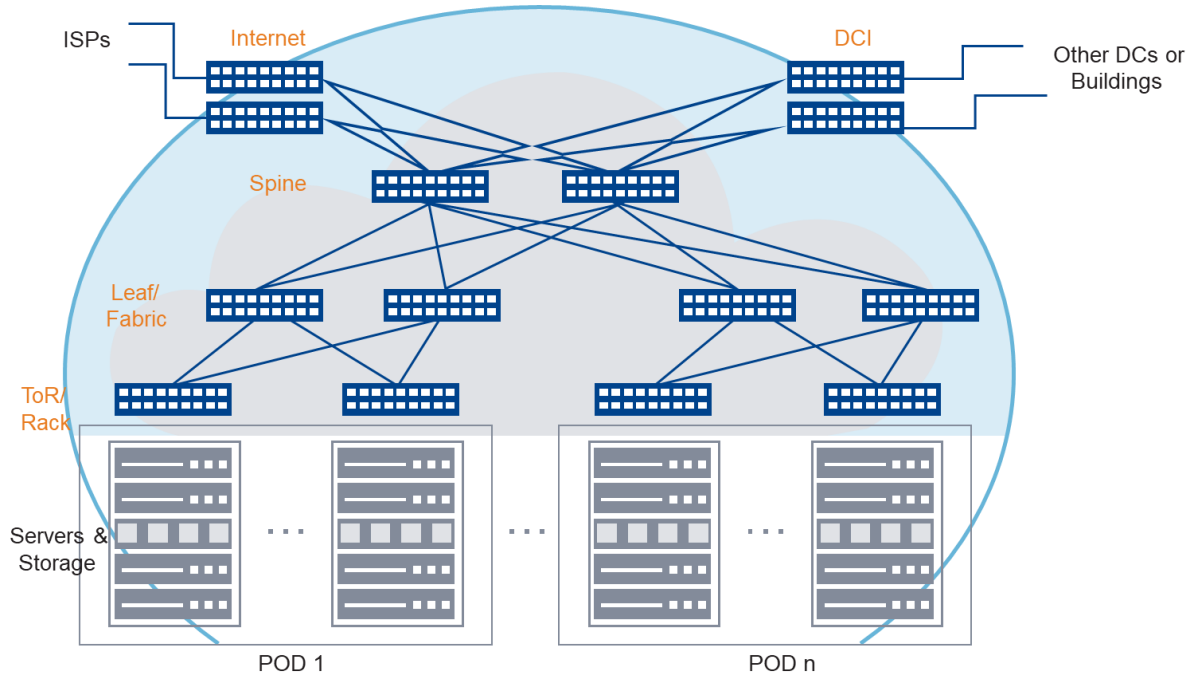

Explosion of Data Traffic in Data Centers

Flattening of Networks for Low Latency

© 2024 Synopsys, Inc.

Main Trends Driving Optics into the Data Center

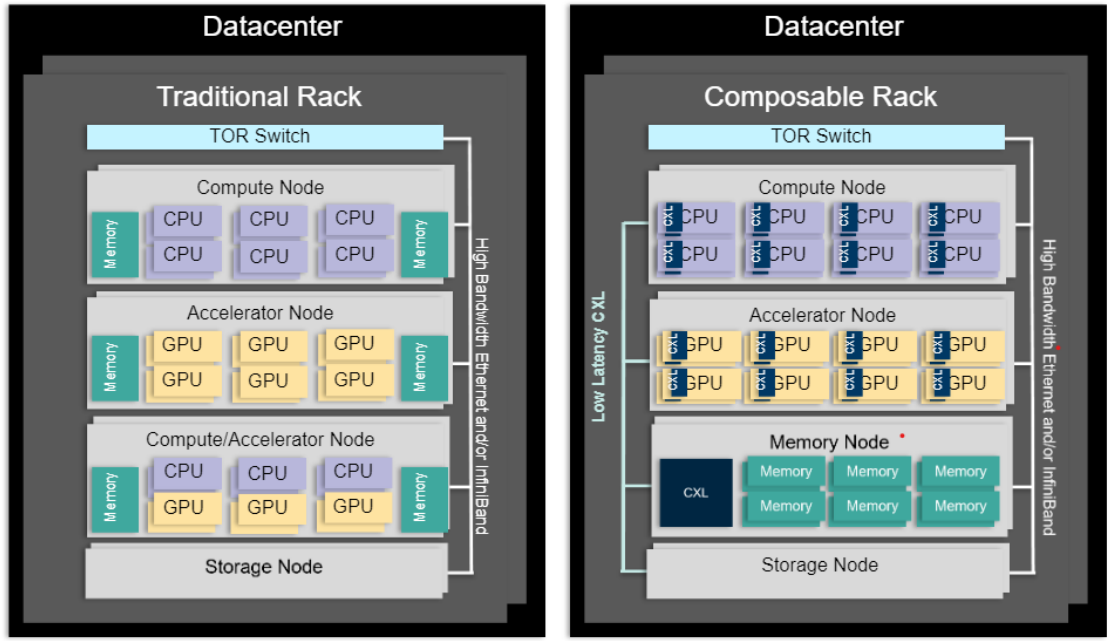

© Hyperion Research 2023

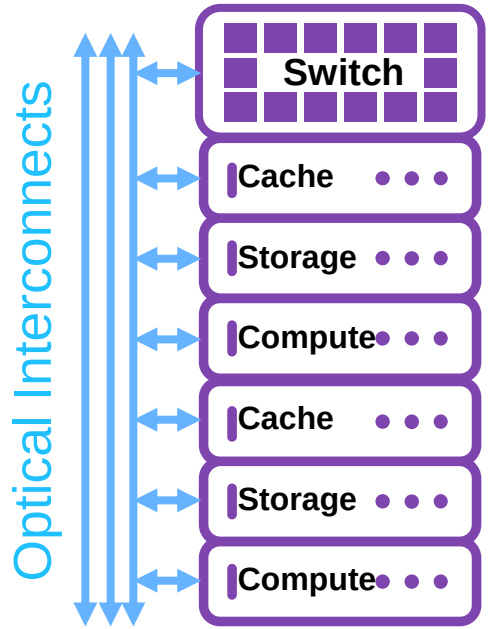

Explosion of Data Traffic in Data Centers

Flattening of Networks for Low Latency

Aggregation of Homogeneous Resources in Rack

© 2024 Synopsys, Inc.

Optical Datacom Moving to Near- & Co-Packaged Optics

Re-thinking Electrical-Optical Interface

- **Complexity** going up
- Proximity of **Electronics and Photonics**
- From single chip to **Multi-Die/Multi-Domain** System
- Power density & Temperature sensitivity
- Reliability
- Changing role for SERDES in the system

  → Impact on **IP**

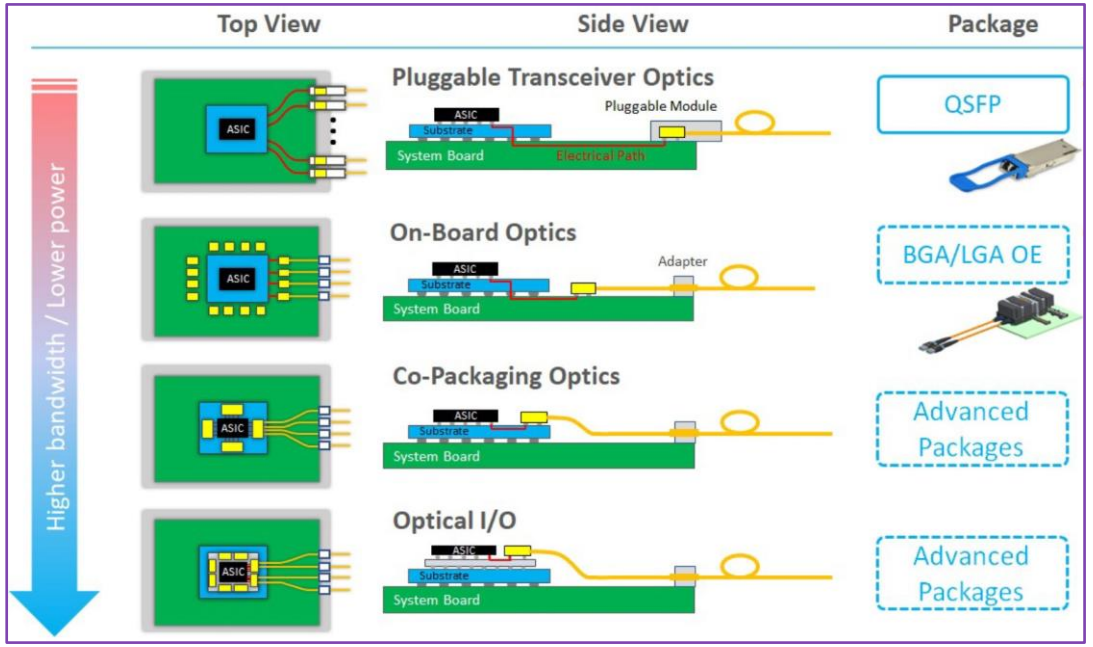

https://ase.aseglobal.com/en/technology/silicon-photonics

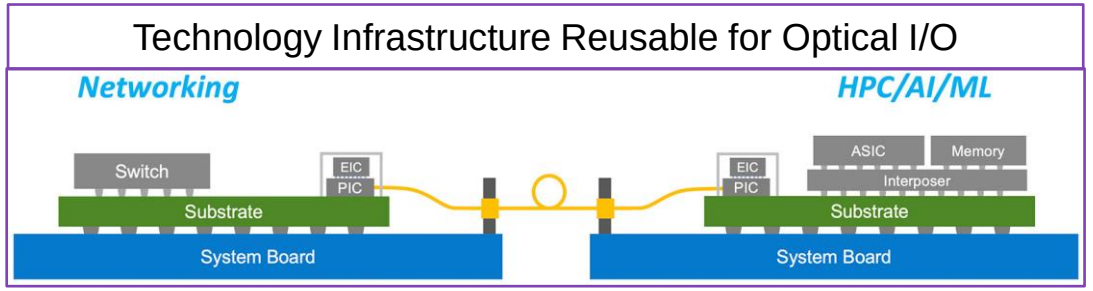

https://ase.aseglobal.com/en/technology/silicon-photonics

© 2024 Synopsys, Inc.

Opportunities for Advances in Electrical-Optical Interface

Accelerating PIC Adoption in Established Markets

© 2024 Synopsys, Inc.

Accelerating PIC Adoption in Established Markets

Re-thinking Electrical-Optical Interface

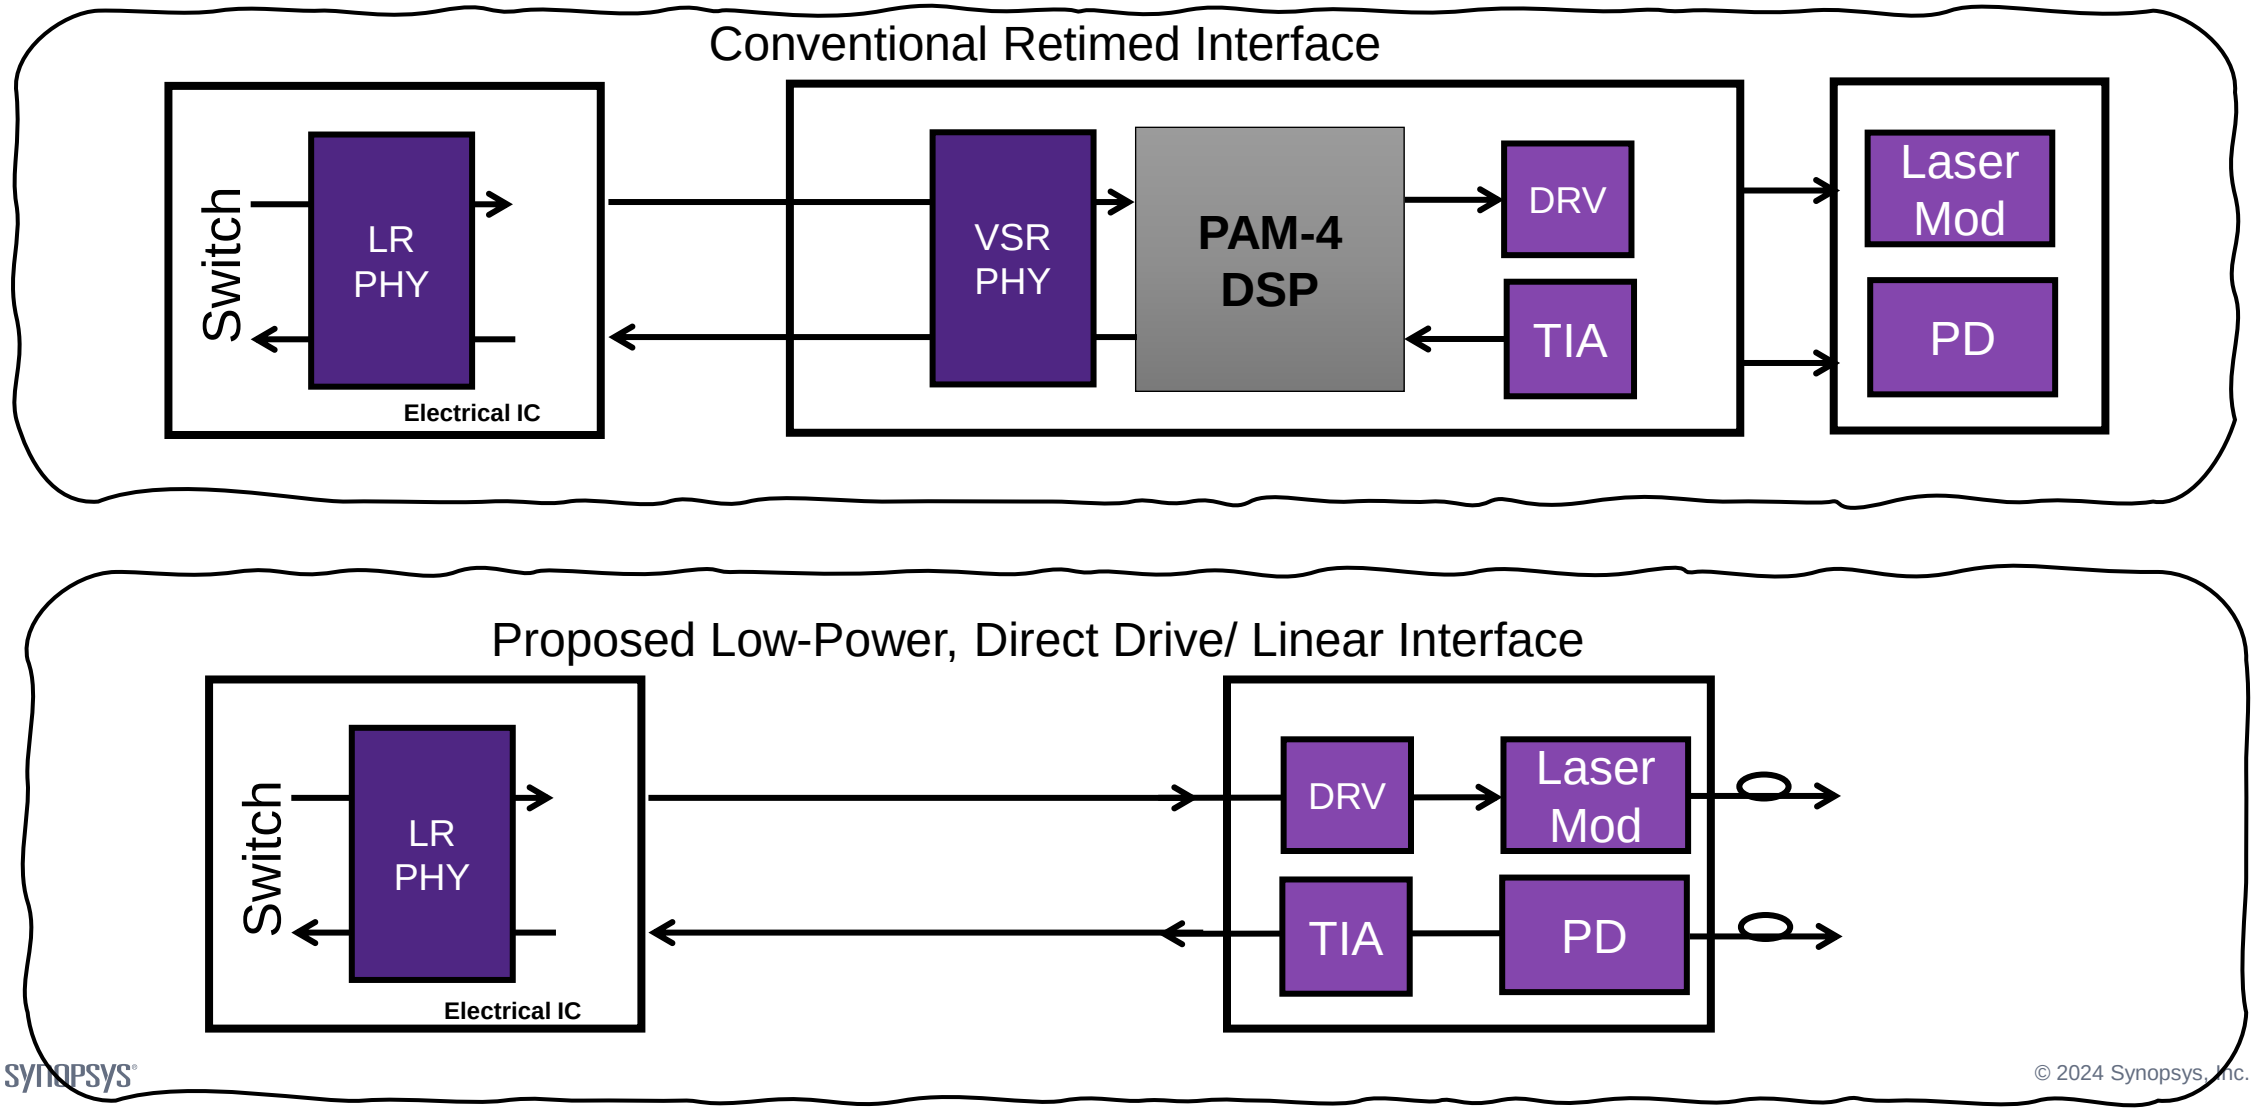

© 2024 Synopsys, Inc.

Next Generation PCIe

For Low Power / Low Latency, Disaggregation with Compute Express Link (CXL)

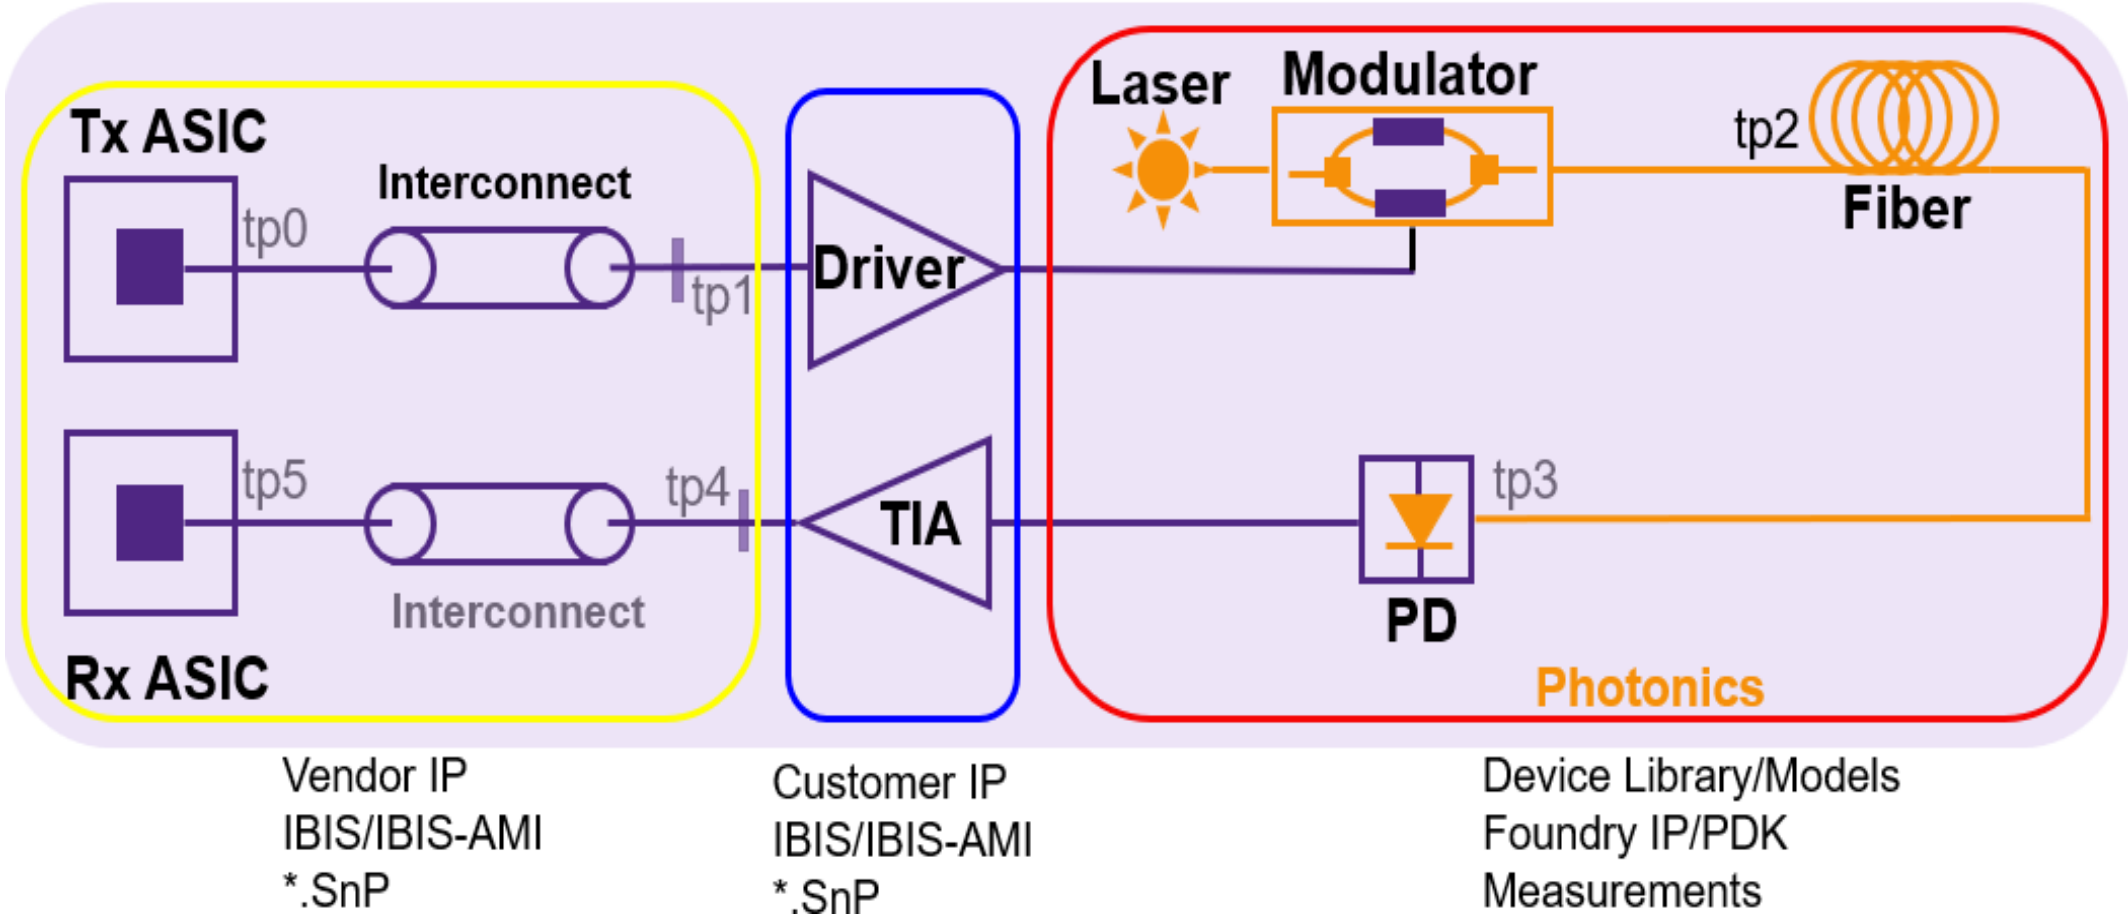

© 2024 Synopsys, Inc.

Photonic Components in an Optical Links

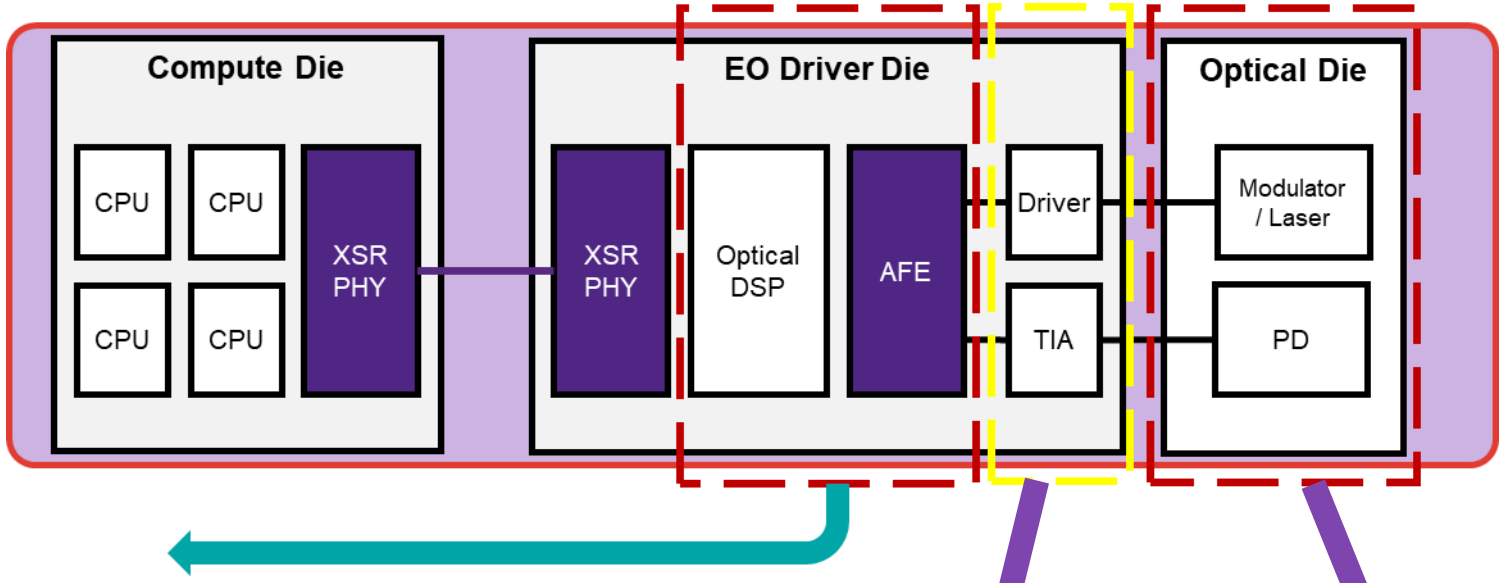

Major Components
- Laser
- Laser Driver
- Modulator (MZM, MRM, EAM)
- Photodiode
- TIA

Fully Electrical
- High TX equalization capability
- High RX equalization capability
- RX May need nonlinear Equalizer

Fully Electrical
- Linear driver with high amplitude
- Low noise, linear TIAs required

Electro-Optical Components
- High BW E-O modulators required
- Laser with Low RIN
- High BW, High sensitivity, and high Responsivity PDs required

© 2024 Synopsys, Inc.

Optical I/O Technology Roadmap

Relentless Innovation

MB/s	GB/s	TB/s		
Optics	Photonics	Photonics & Electronics	Embedded Optics	E/O SiP
Discrete components	Integrated Photonics	Silicon Photonics	Co-Packaged Optics	xPU/D2D Optical I/O
Pre-2020		2020 .. 2023	Future	

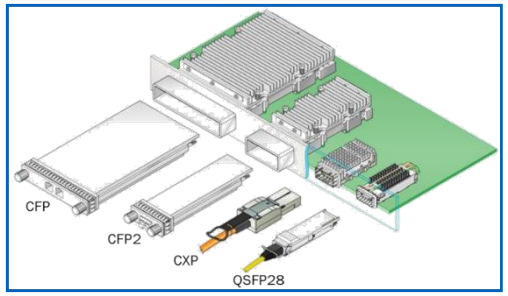

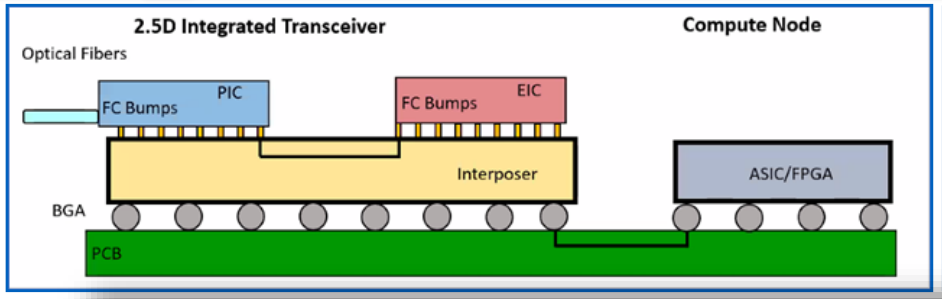

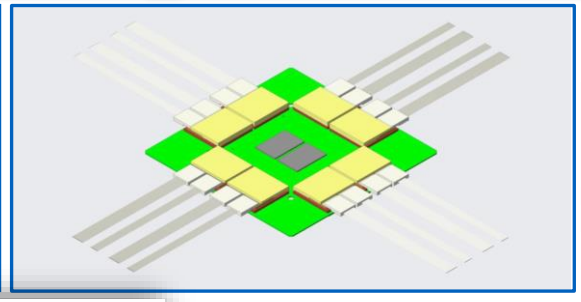

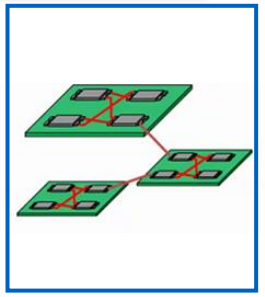

Source: OIF

PCI-SIG® Exploring an Optical Interconnect to Enable Higher PCIe Technology Performance

PCI-SIG establishes a new workgroup to seek industry feedback in developing an optical interconnect

August 02, 2023 11:00 AM Eastern Daylight Time

BEAVERTON, Ore.--(BUSINESS WIRE)--PCI-SIG® today announced the formation of a new workgroup to deliver PCI Express® (PCIe®) technology over optical connections. The PCI-SIG Optical Workgroup intends to be optical technology-agnostic, supporting a wide range of optical technologies, while potentially developing technology-specific form factors.

"We have seen strong interest from the industry to broaden the reach of the established, multi-generational and power-efficient PCIe technology standard by enabling optical connections between applications"

Tweet this

"Optical connections will be an important advancement for PCIe architecture as they will allow for higher performance, lower power consumption, extended reach and reduced latency," said Nathan Brookwood, Research Fellow at Insight 64. "Many data-demanding markets and applications such as Cloud and Quantum Computing, Hyperscale Data Centers and High-Performance Computing will benefit from PCIe architecture leveraging optical connections."

© 2024 Synopsys, Inc.

Opportunities for PIC due to the Pressure on PCIe/UCIe to Deliver

Accelerating PIC Adoption in Established Markets

© 2024 Synopsys, Inc.

More Standards Are Adopting Optics

OIF is Leading the Move; PCI-SIG and UCIe Consortium are Joining

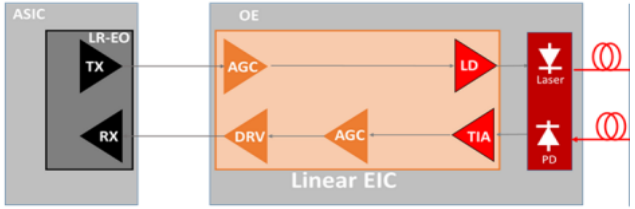

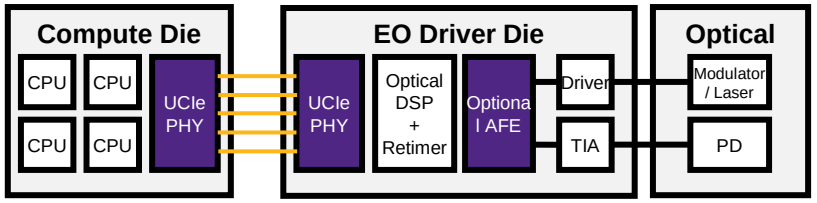

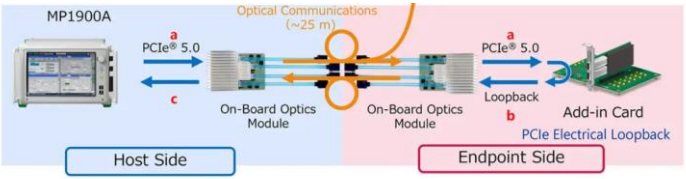

- 112G Linear Optics
- 56G Linear Optics
- 800G ZR
- 800G LR
- New Energy efficient optical interfaces

- Defined UCIe Retimers for off Package Optical interconnect

- Discussions on Optical interface for PCIe interconnects – Retimed or Un-retimed
- New workgroup to deliver PCIe technology over optical connection

© 2024 Synopsys, Inc.

CPO with D2D UCIe

For Low Power / Low Latency, Dense Networks Within DC

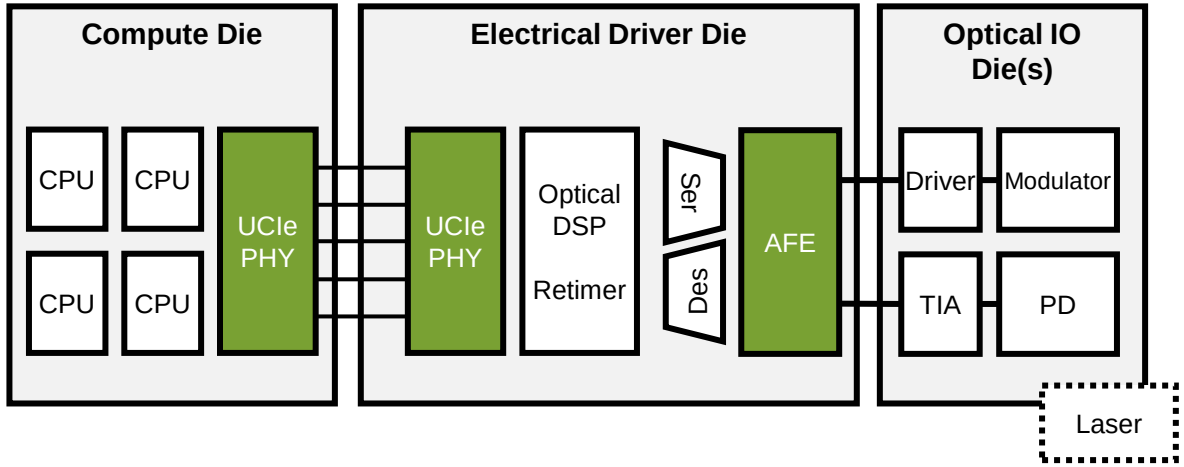

- A three-die implementation with UCIe PHY allows true open multi vendor integration
- Laser integration is shown as an option
- Robust interoperability at the cost of slightly higher power

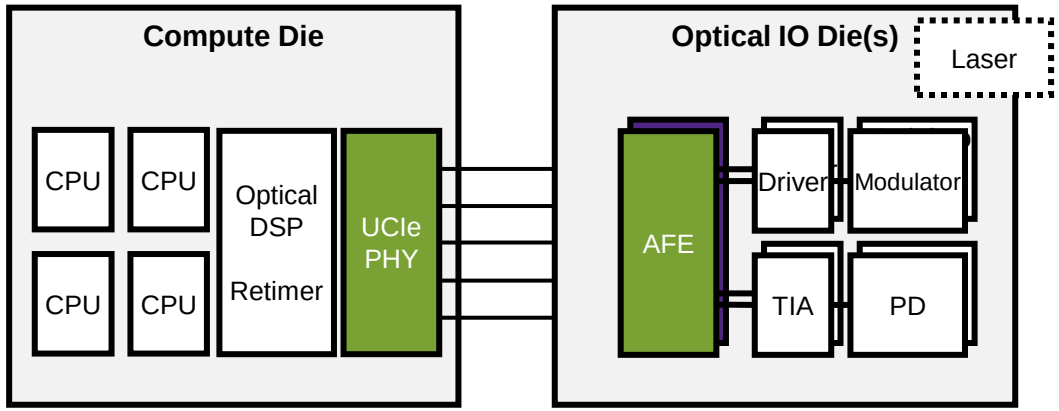

- A two-die implementation with UCIe PHY in allows custom implementation for a reach and vendor
- Laser integration is shown as an option
- Saves power as avoids a stage of interconnect

© 2024 Synopsys, Inc.

Opportunities for the EDA Industry

Due to Accelerating PIC Adoption and 3D-HI

© 2024 Synopsys, Inc.

Techonomic Disruptors in SysMoore Era

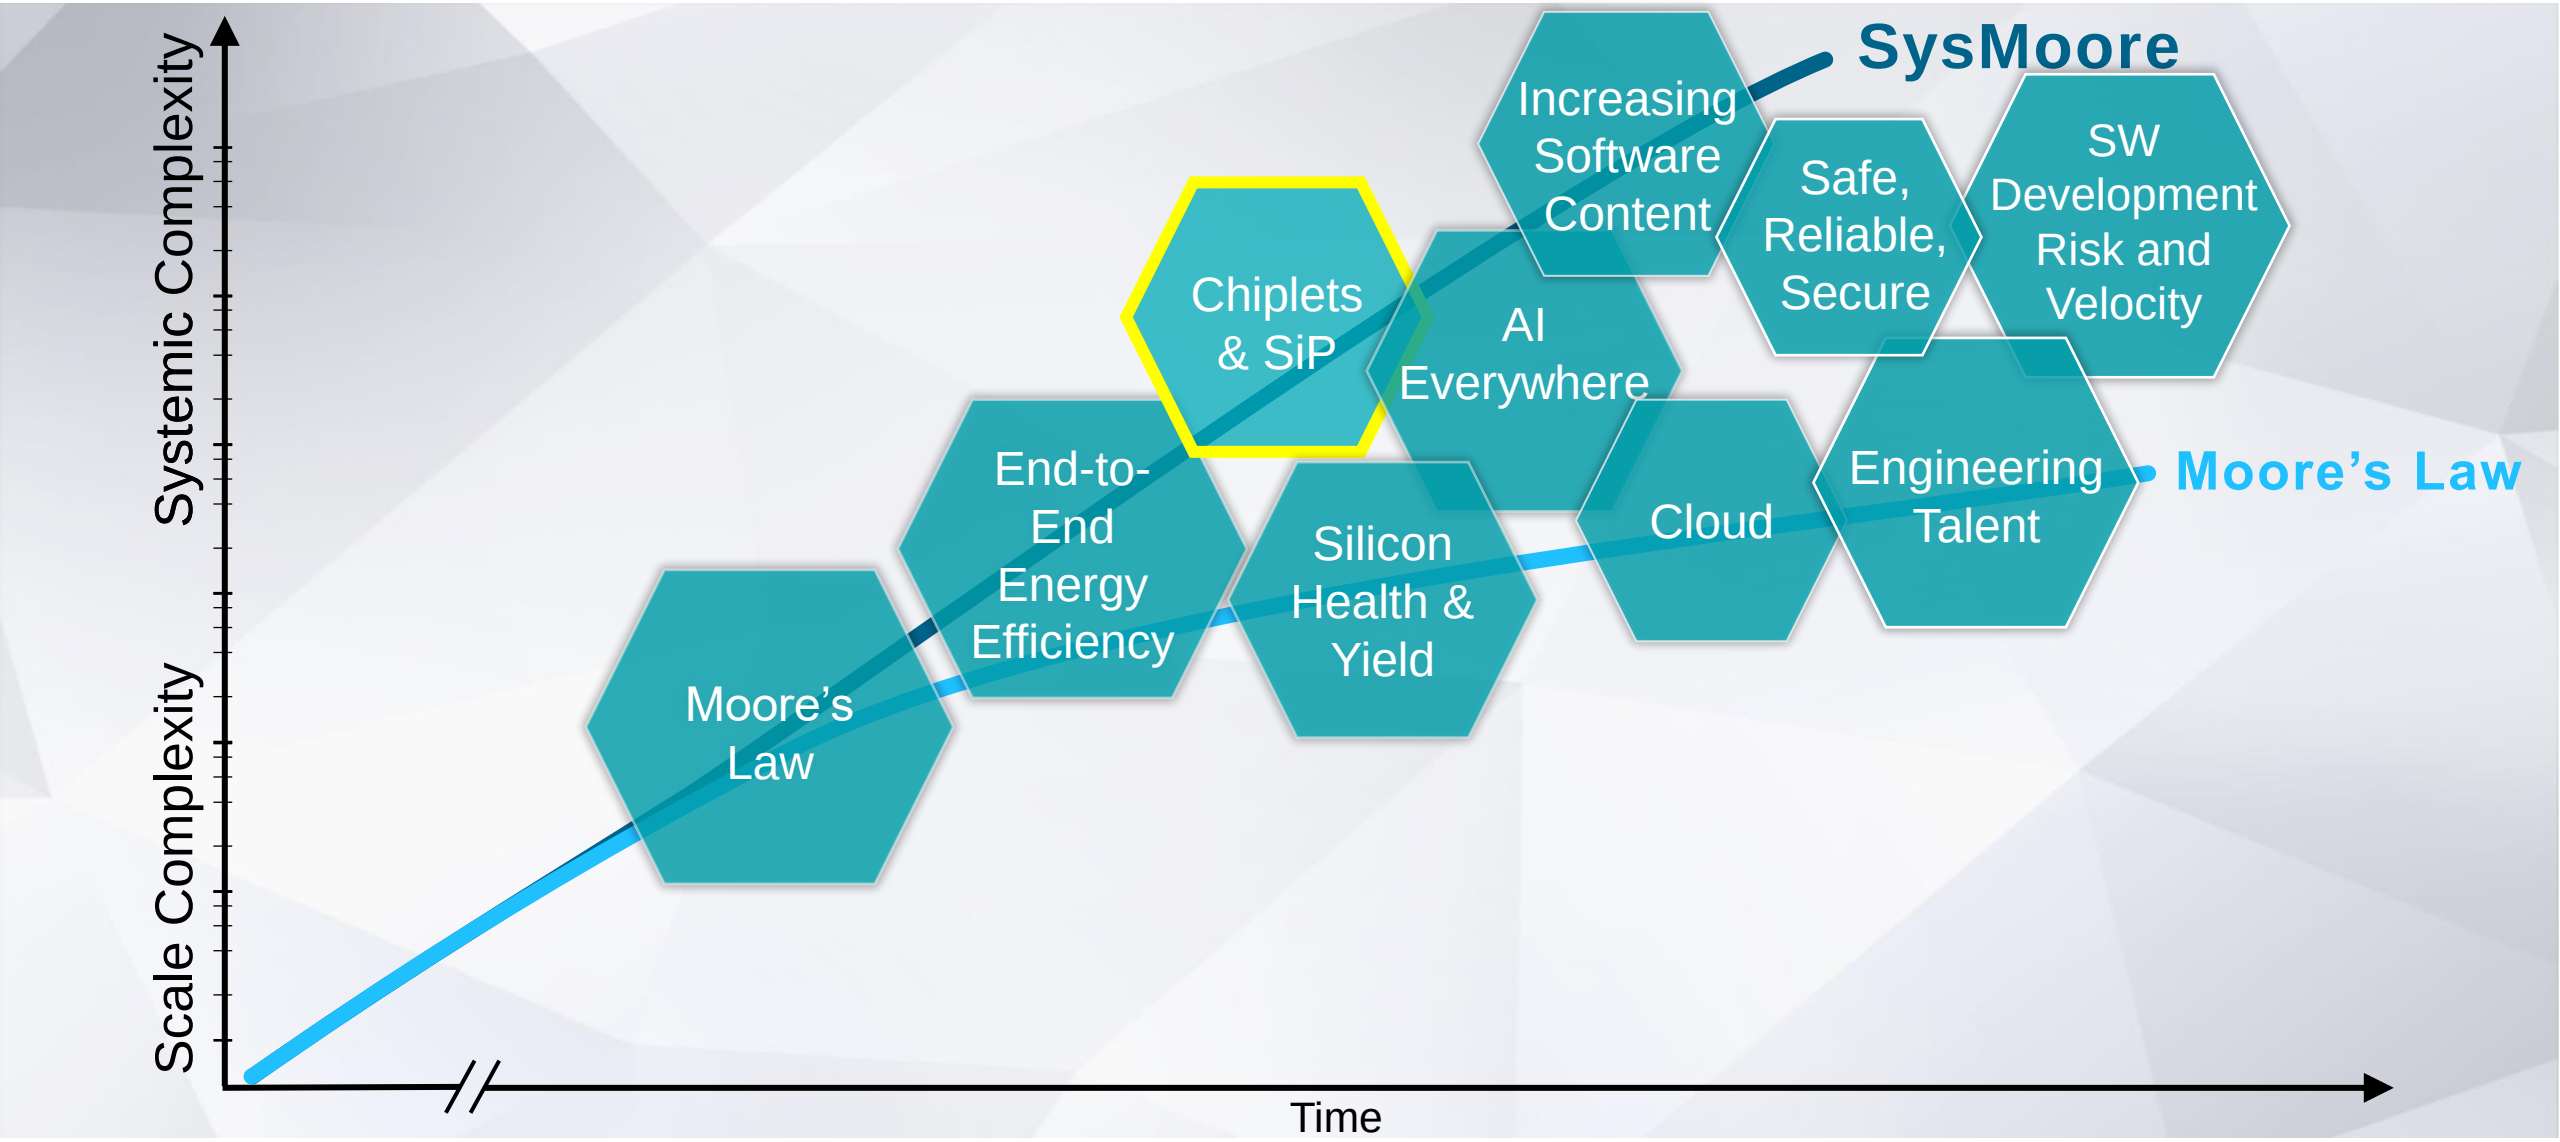

© 2024 Synopsys, Inc.

Multi-Die Systems are Here

AMD
Epyc Server CPU

- Scale & Disaggregate IO
  - 8 Compute Die +
  - 1 I/O Die
- Substrate Pkg

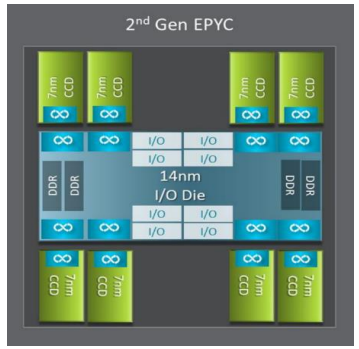

Apple
M1 UltraMax CPU

- Scale use case
  - 2 Compute Die
- Silicon Interposer Pkg

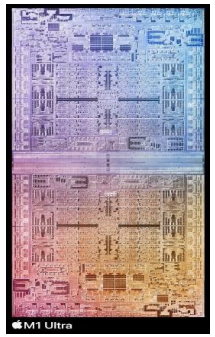

Tesla
Dojo Training Tile

- 25 chiplets packaged in TSMC's InFO with 6 RDL (3 thick for power)

Amazon
Graviton3 Server CPU

- Disaggregate IO
  - 1 Compute Die +
  - 4 DDR Die + 2 I/O Die
- Substrate Pkg

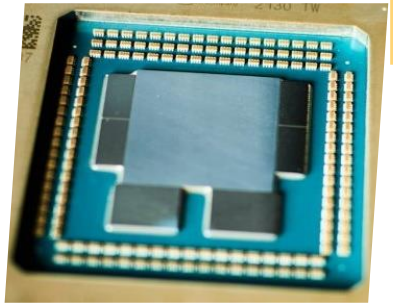

Intel
Ponte Vecchio Xe-HPC GPU

- Scale & Disaggregate IO
  - 16 Compute Die +
  - 8 Cache Die + 2 Base Die
  - 8 HBM Die + 2 IO Die
- EMIB + Foveros Pkg

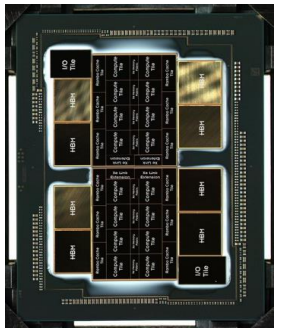

© 2024 Synopsys, Inc.

Multi-Die System Design Adds New Complexities

Opportunities for the EDA Industry

Co-design

- Leveraging EDA tools
- Early model-based architectural exploration

Reliable Die-to-Die Connectivity

- Include features to meet multi-die system requirements

Designing at the system level

- System-level validation
- System sign off
- Silicon health during entire operation

Thermal challenges

- Good thermal and electrical modeling
- Performance per watt becomes critical

Test challenges

- Known good die (KGD) is required

Source: TechSearch International

© 2024 Synopsys, Inc.

Multi-Die / Multi-Domain System-of-Chips Design Flow

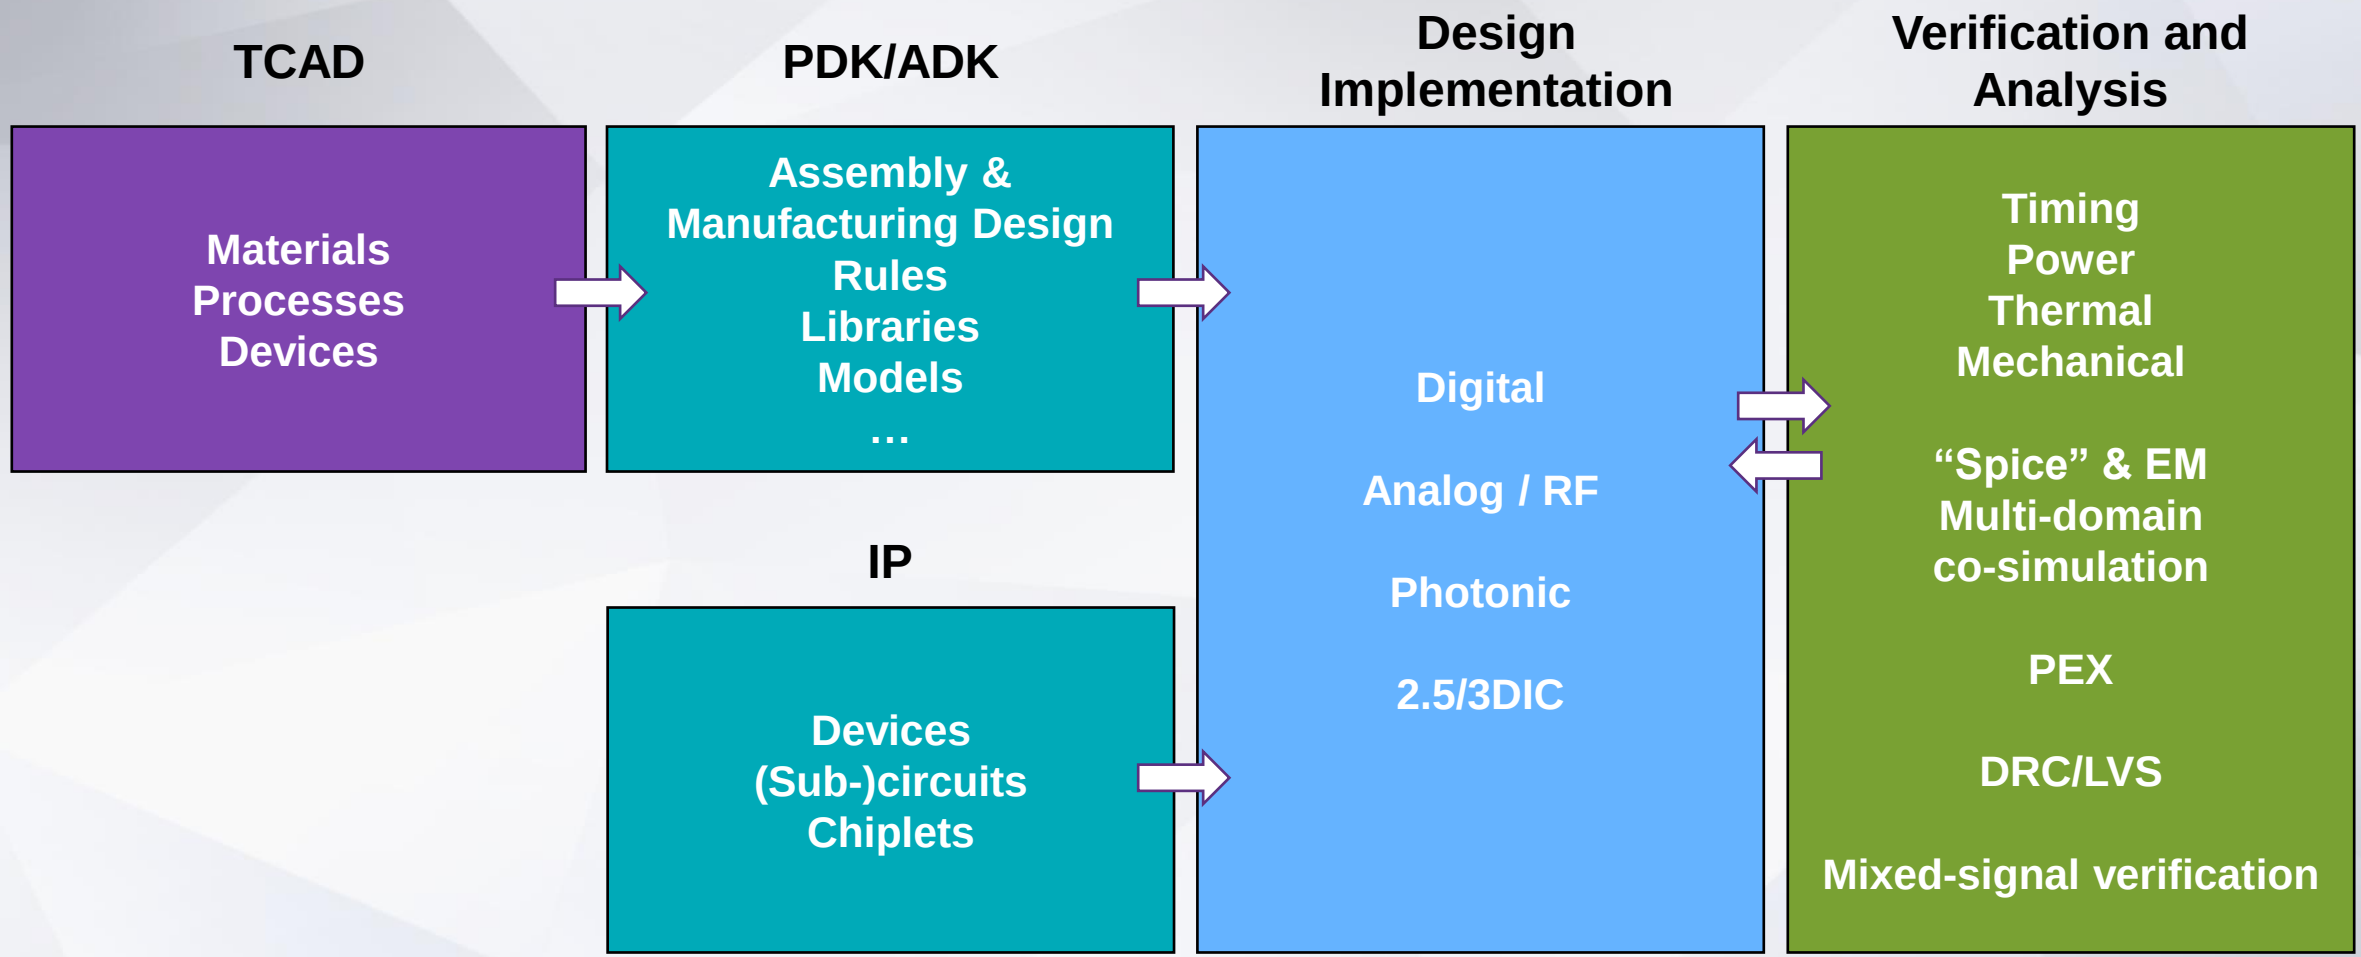

© 2024 Synopsys, Inc.

Multi-Die System Design

Critical Design Aspects in 3D Heterogenous Integration

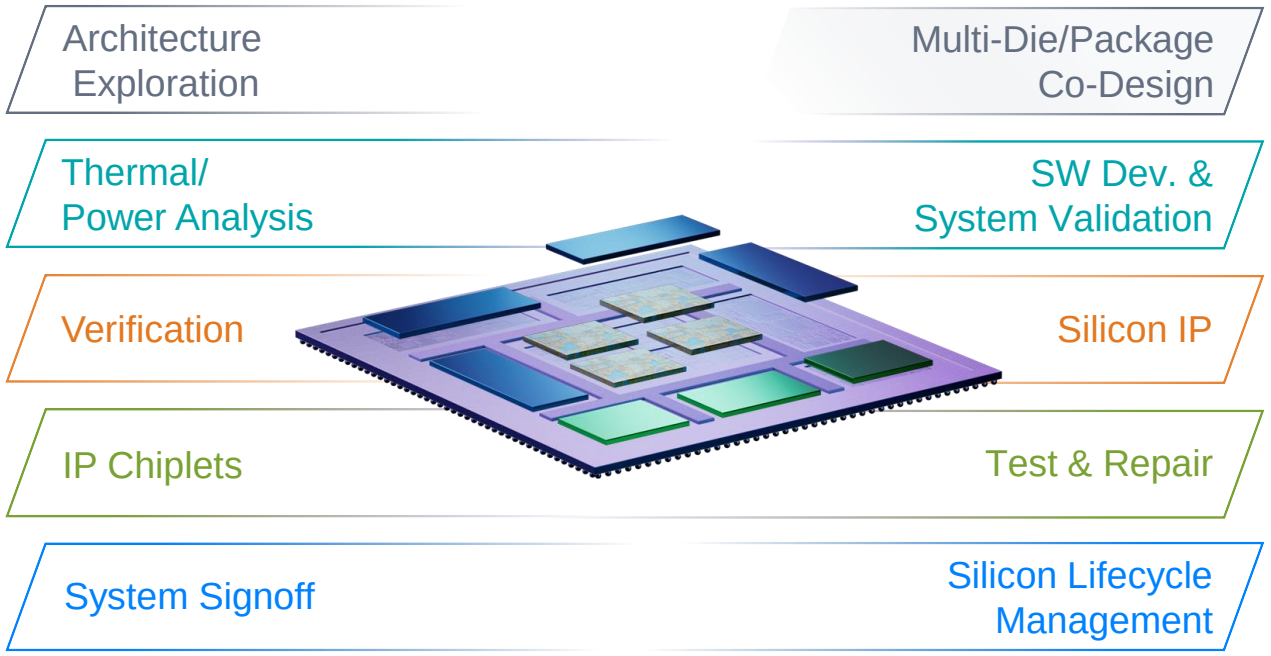

Architecture Exploration	Optimize thermal, power, and performance with early exploration and partitioning
Software Dev. & Validation	Rapid software development and validation with high-capacity emulation & prototyping
Design Implementation	Efficient die/package co-design with unified exploration-to-signoff platform and robust IP
Manufacturing & Reliability	Improve health, security and reliability with holistic test and lifecycle management solutions

© 2024 Synopsys, Inc.

Thank You


SYNOPSYS®