Essence of High Speed Connectivity: When Big Data Meets Big Compute

Wendy Chen, APAC IP Sales and Eco-system Group Director
D&R IP SoC DAYS
Shanghai,
Sep 8th, 2018

cadence®

The Cyber World is Growing BIG and FAST

Facebook
8 Billion video views daily

Amazon
6,000 orders
closed each minutes

YouTube
300 hours of video
uploaded each minute

NETFLIX
140 Million hours of
Content watched daily

WeChat
38 Billion daily messages

© 2018 Cadence Design Systems, Inc. All rights reserved.

Evolution of Data Processing: Mainframe/Client/Cloud/Accelerator

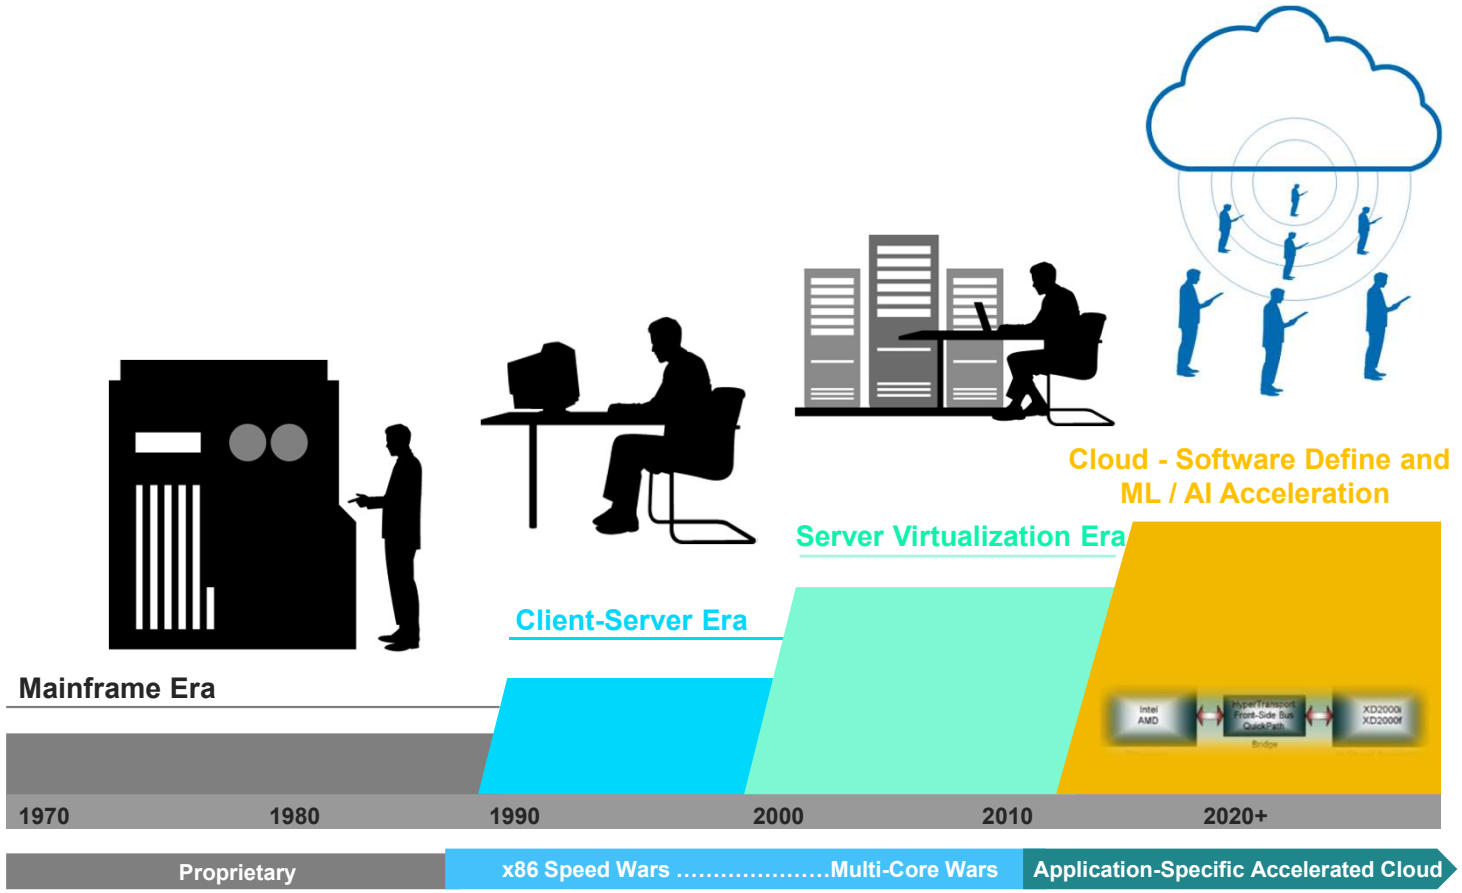

© 2018 Cadence Design Systems, Inc. All rights reserved.

What Is Happening Today in the Enterprise World?

Per IDC: In 2018, 75% of enterprise developments will include AI / ML / DL

Healthcare and Life Sciences

- Molecular dynamics
- Drug interactions

Financial Services

- Risk analysis
- Fraud detection

Government

- Disaster planning
- Emergency service allocation

Manufacturing

- Smart facilities
- Predictive maintenance

Retail

- Consumer analytics
- Product recognition

Energy

- Production management
- Environmental impact

Transportation

- Driver assistance
- Autonomous driving

Travel

- Traveler recognition
- Intelligent overbooking

© 2018 Cadence Design Systems, Inc. All rights reserved.

Why: Data Is the New Oil!

AI can make the smart grid smarter	**20%**	Boost in energy production
AI can make health diagnoses faster, for better treatment and improved health care	**40%**	Higher chances of finding rare genetic disorders
AI can make retailers know shoppers pattern, to get them to come back again	**65%**	Reduction in lost sales due to logistics and delivery, accuracy in inventory management
AI can make bionic hands know what they are grabbing	**10X**	Faster than contemporary prosthetics

© 2018 Cadence Design Systems, Inc. All rights reserved.

Datacenter Opportunity: ~$18B SAM in 2020

Workload-optimized, high-performance compute, connectivity, accelerators – ML/DL/AI

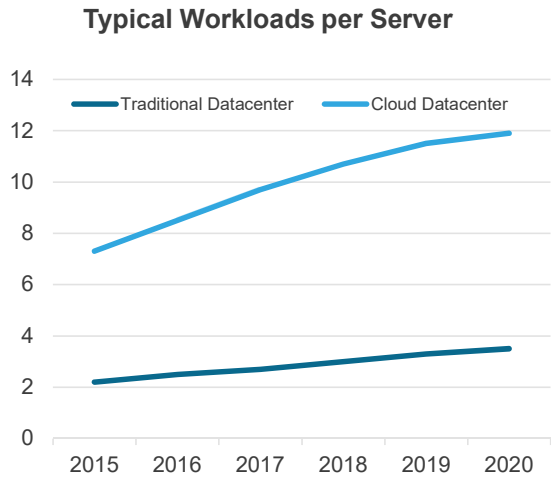

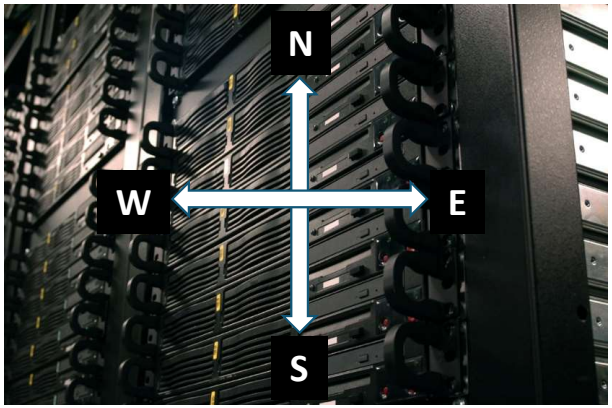

Hyperscale Optimization CPU

- Workload optimized
- Machine learning
- Deep learning
- Accelerator offloads

Rack-Level Connectivity

- Leaf /spine
- Memory pool (HBM)
- Connectivity / SiP
- Reduced latency
- Mesh / 3D-torus / fabric

Scale Out Clusters

- DNN
- SSD / NVMe
- Coherency
- VM / containers
- Mesh/3D-torus

© 2018 Cadence Design Systems, Inc. All rights reserved.

Market Trend – Datacenter Cloud Connectivity

Majority of connectivity will be 50G/100G by 2020

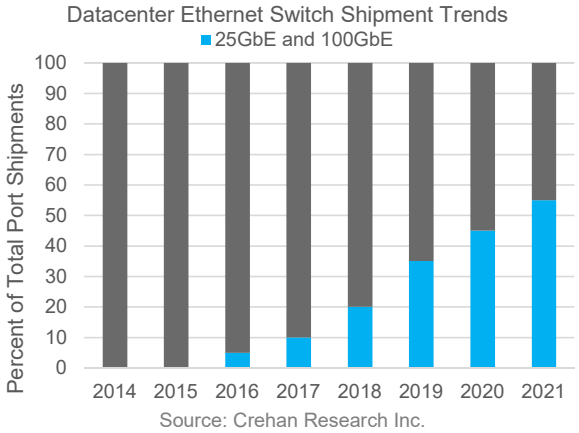

1GbE: Single to Multiple Port 1GbE

10GbE: Single to Multiple Port 10GbE

40GbE: Single 40GbE or Quad-Port 10GbE

25GbE: Single-Port 25GbE

50GbE: Dual-Port 25GbE or Single-Port 50GbE

100GbE: Single-Port 100GbE or Quad-Port 25GbE

Segments	Speed of Adoption	2016	2017	2018	2019	2020
Tier 1 Cloud Datacenter (>1M Servers)	Early Adopter	10GbE > 40GbE	40GbE > 50GbE	50GbE > 100GbE	100GbE	100GbE+
	Majority Adopter	10GbE	10GbE > 25GbE	25GbE > 50GbE	50GbE > 100GbE	50GbE > 100GbE
Tier 2/3 Cloud Datacenter	Early Adopter	1GbE > 10 GbE	10GbE > 25GbE	25GbE > 50GbE	50GbE > 100GbE	50GbE > 100GbE
	Majority Adopter	1GbE > 10 GbE	10 GbE	10GbE > 25GbE	25GbE > 50GbE	25GbE > 50GbE
Enterprise / On Premises	Early Adopter	1GbE > 10 GbE	10GbE > 40GbE	10GbE > 40GbE/50GbE	10GbE/40GbE > 50GbE	50GbE
	Majority Adopter	1GbE	1GbE > 10GbE	1GbE > 10GbE	10 GbE	10 GbE

Source: Dell'Oro Group 2017, delloro.com

© 2018 Cadence Design Systems, Inc. All rights reserved.

cadence

Artificial Intelligence Needs the Most Intelligent Interconnect IP

Intelligent Offload

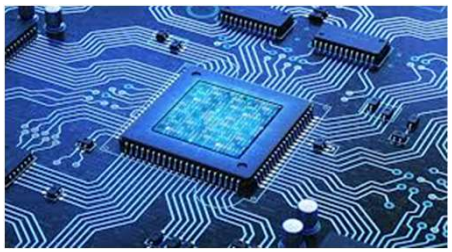

SoC – Accelerators / Chiplets

Scalability / Efficiency

Higher Connectivity Speeds

Faster Data Processing

Better Data Security

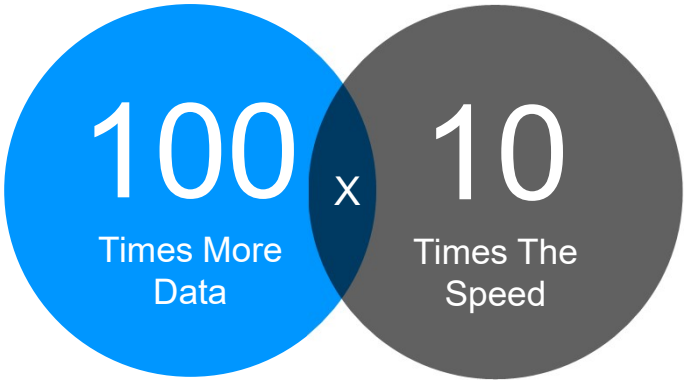

© 2018 Cadence Design Systems, Inc. All rights reserved.

cadence

What's the Primary Goal (How): Silicon Design Specific to ML/AI

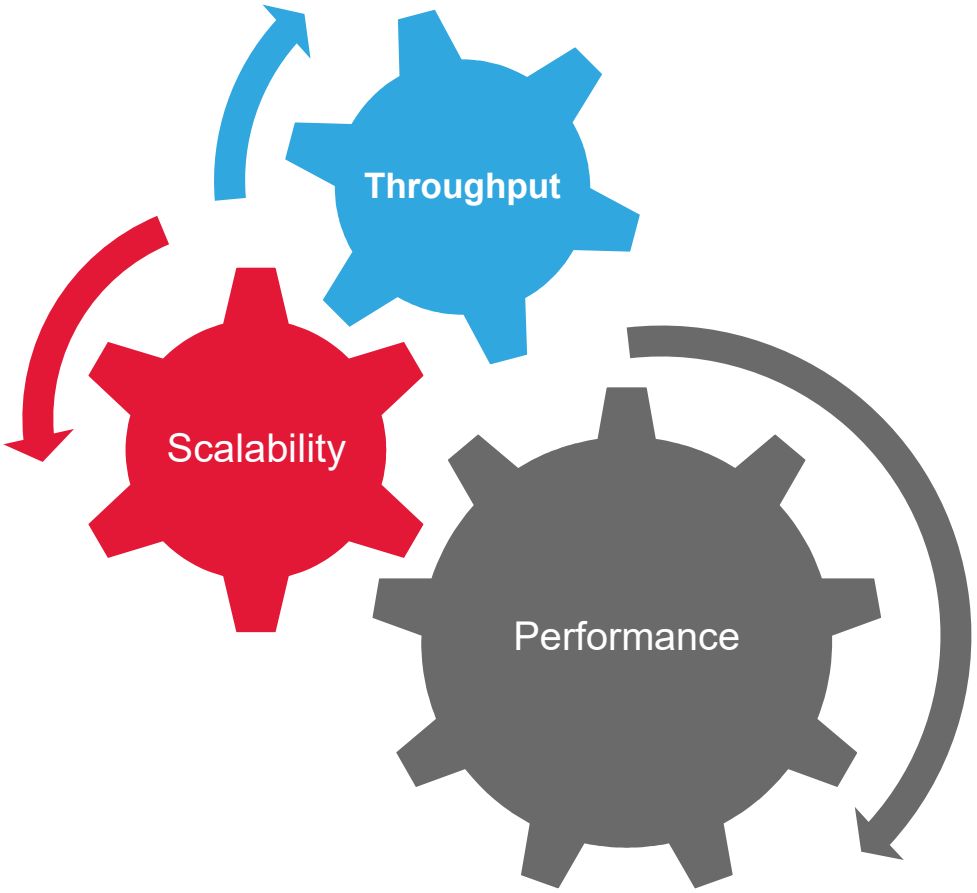

- Run many jobs in parallel
- Eliminate queue wait times
- Scalable clusters to meet the demand (clients)
- Optimize the resources for appropriate task
- Optimize performance for specific task

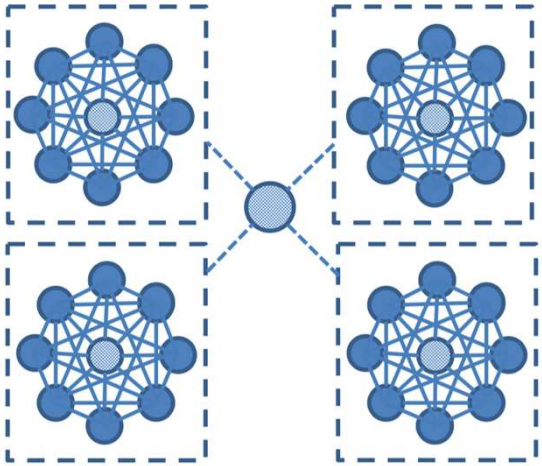

© 2018 Cadence Design Systems, Inc. All rights reserved.

Machine Learning Market Landscape

Machine Learning Market Size

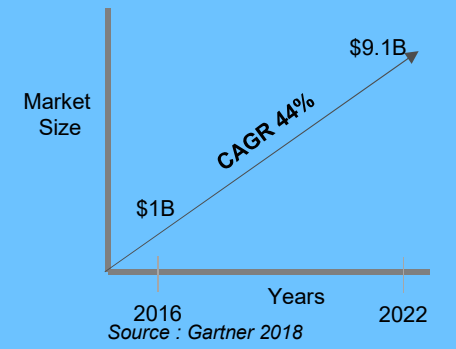

Source : Gartner 2018

Machine Learning by Application

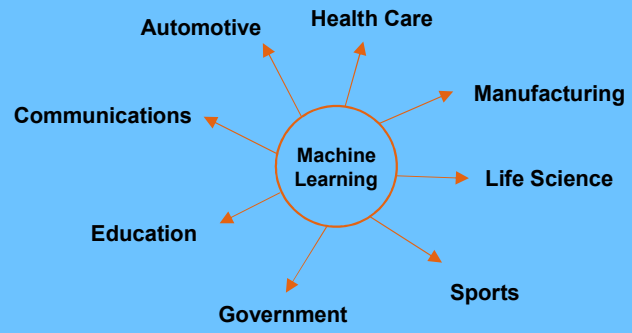

Types of Machine Learning

- Supervised learning [task driven]
- Unsupervised learning [data driven]
- Reinforcement learning [adaptive]

© 2018 Cadence Design Systems, Inc. All rights reserved.

Deep Learning Semiconductor Forecast

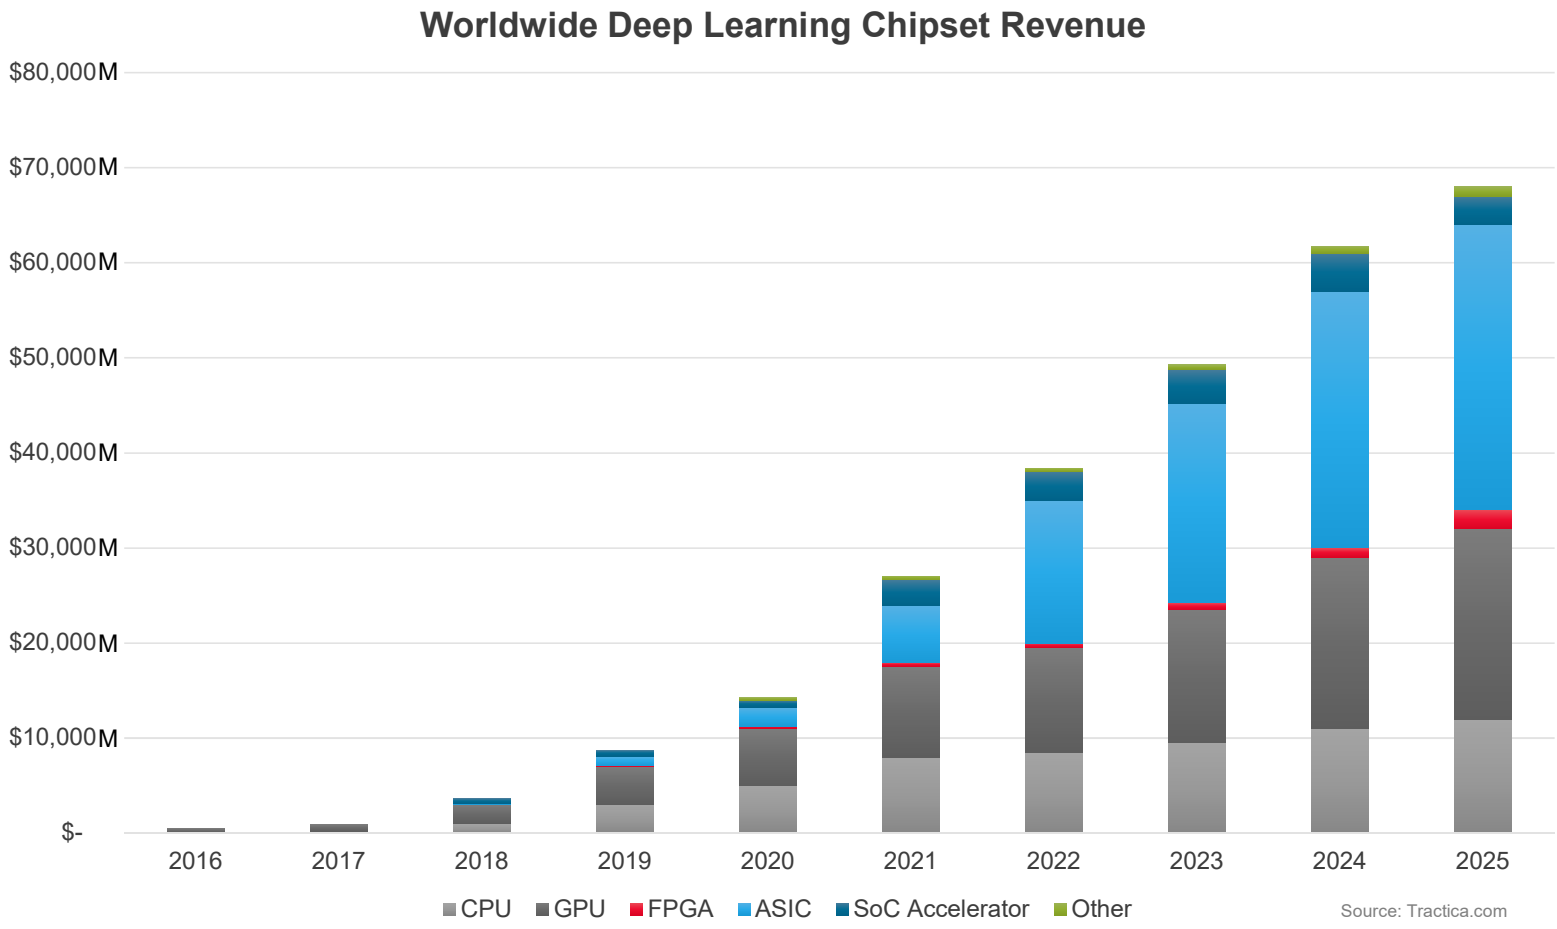

© 2018 Cadence Design Systems, Inc. All rights reserved.

cadence

Deep Learning Architecture Example – Intel

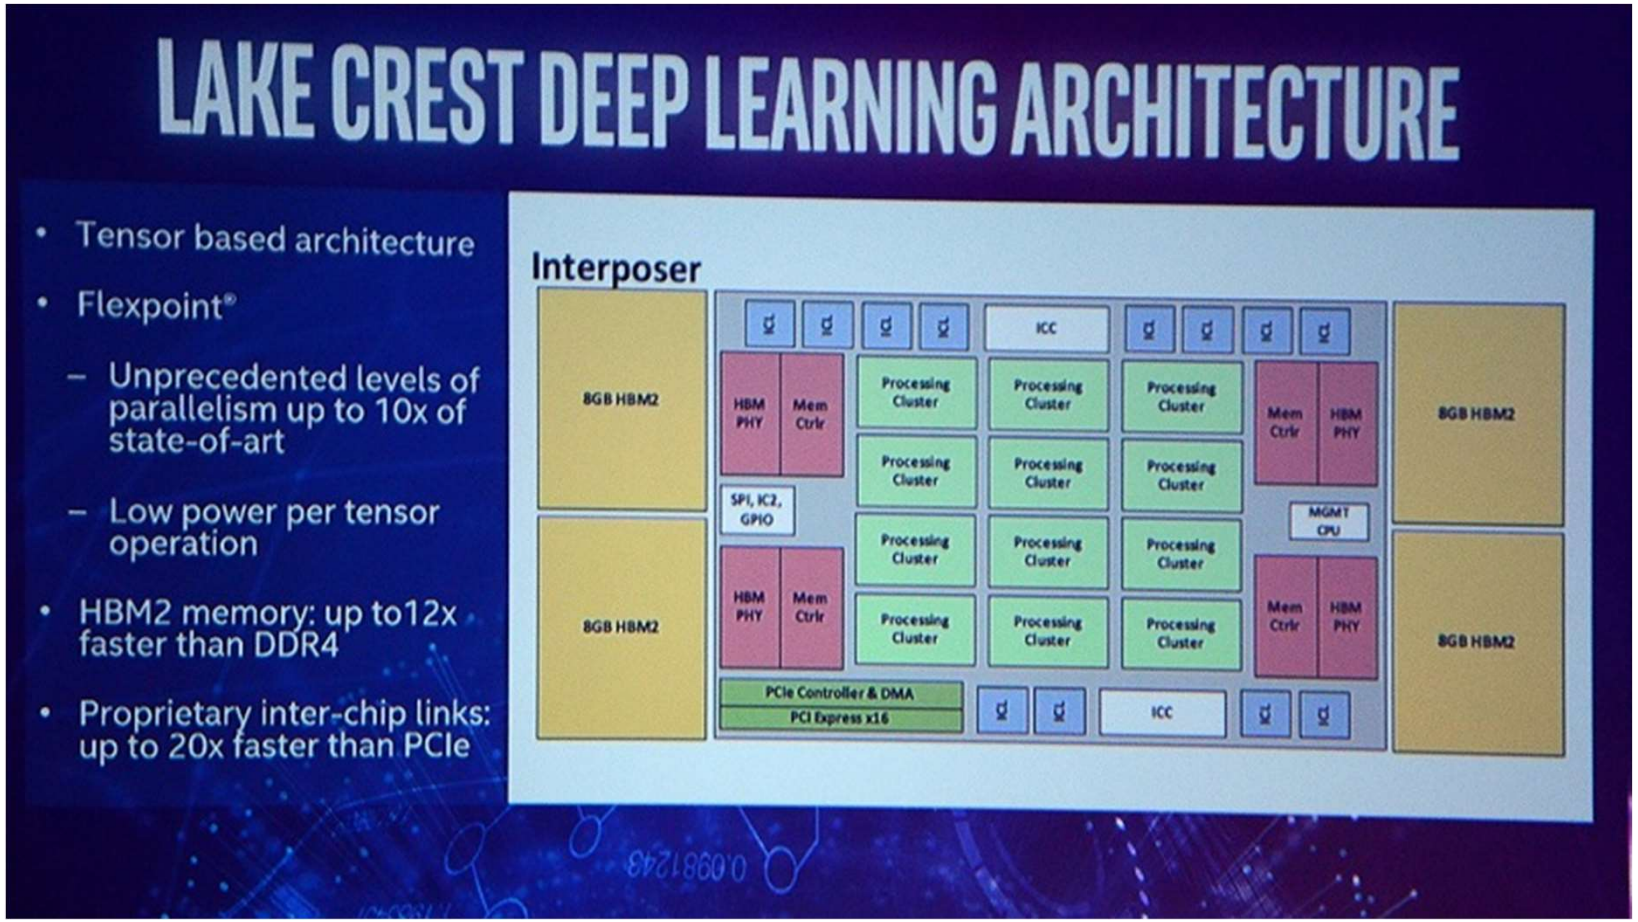

Source : https://www.eteknix.com/intel-lake-crest-ai-accelerators/

© 2018 Cadence Design Systems, Inc. All rights reserved.

Drivers for AL/ML/DL Silicon Design

- AI applications are driving the development of new silicon & system architecture
- Process Density & Power key decision driver for silicon design
- Key Trends Influencing AI Silicon Architecture :
  - Processor Architecture
  - Memory Bandwidth
    - Multiple Options : On-Chip , HBM , GDDR
- PCI Express® (PCIe®) Gen 5 / CCIX / GenZ fabric connectivity
- IO Connectivity 25G / 56G /112G
- Die-Die connectivity within package – scalability
- Silicon Photonics

© 2018 Cadence Design Systems, Inc. All rights reserved.

IP Requirements for Building AI / ML Silicon

- Connectivity
  - 25G/32G/56G/112G
  - 32G PCIe® Gen5
- Memory
  - HBM
  - GDDR6
- Inter-die connectivity
  - D2D
  - 112G XSR

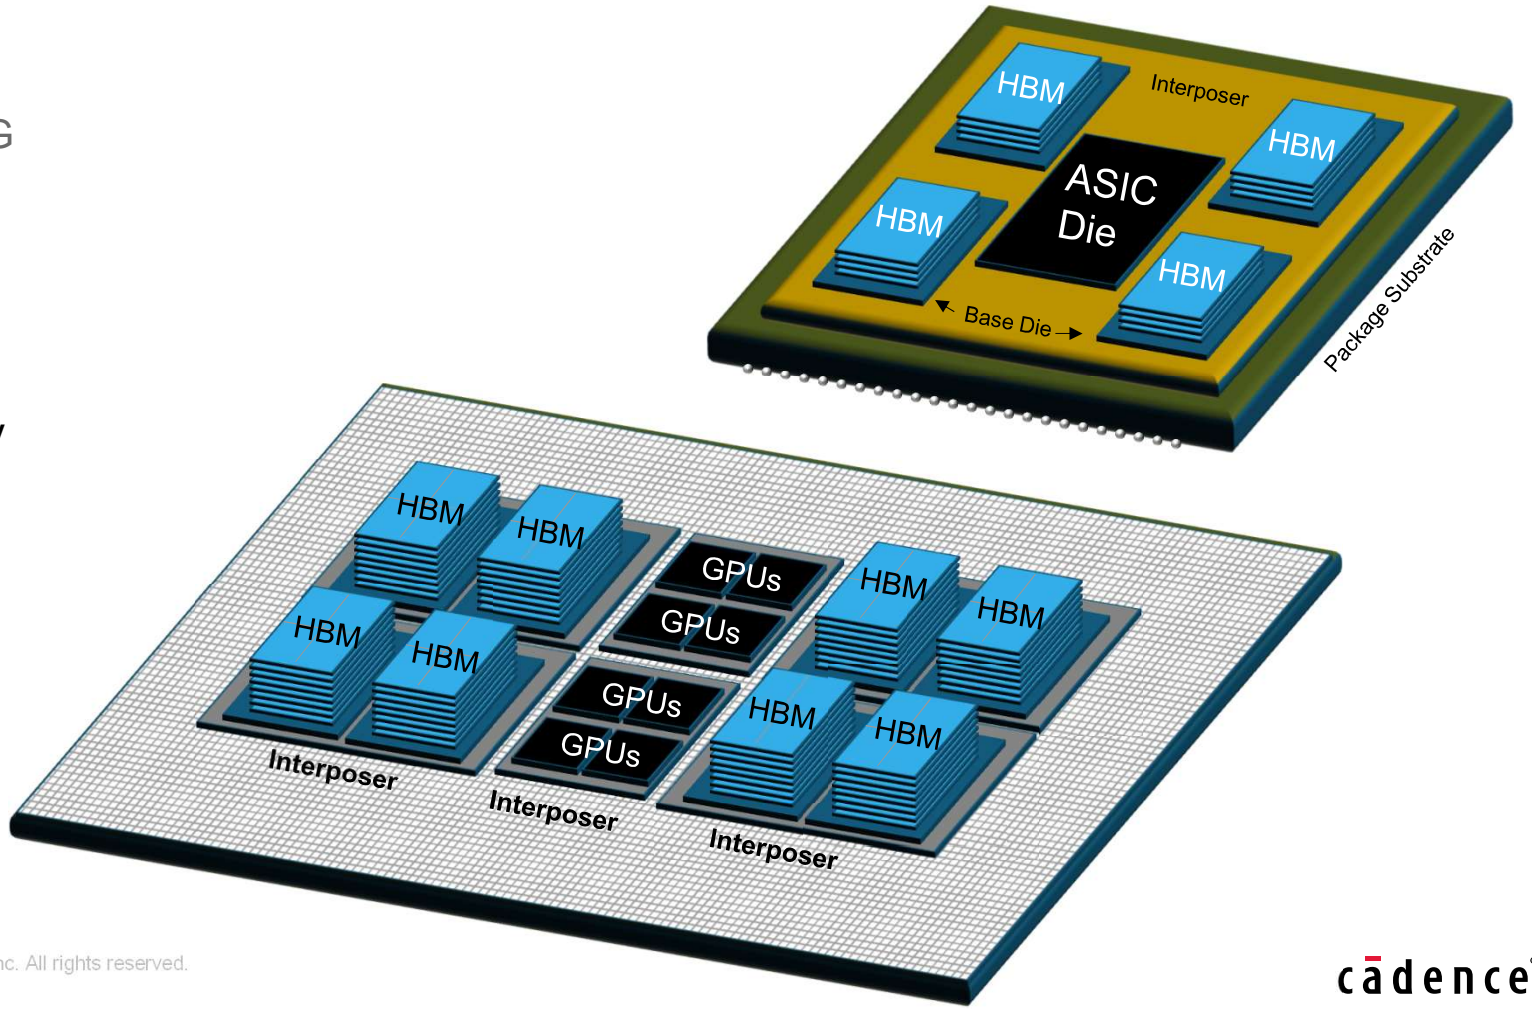

© 2018 Cadence Design Systems, Inc. All rights reserved.
cadence

Challenges at 7nm and Beyond

Need for system in package solutions

Source: AMD, IEDM 2017

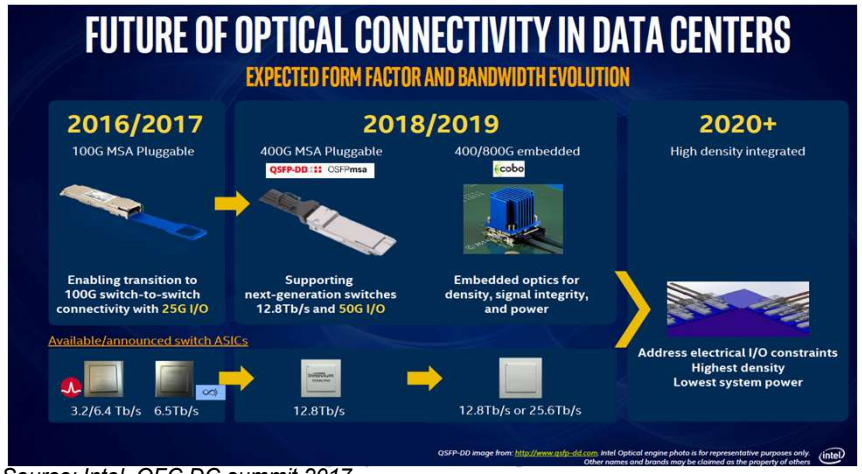

Source: Intel, OFC DC summit 2017

- New breed of data-intensive vision, graphics, AI, ML, and accelerator applications driving need for increased per socket compute power, increasing monolithic die cost and manufacturability challenges

Source: AMD, IEDM 2017

- Hyper-scale datacenters, driving SoC bandwidth, size and need for hybrid integration (analog, optical, accelerators, etc.)

cadence

Die2Die Interconnect
Overview and applications

- Rising monolithic die cost and limitations of on-die integration driving move to die disaggregation and SiP

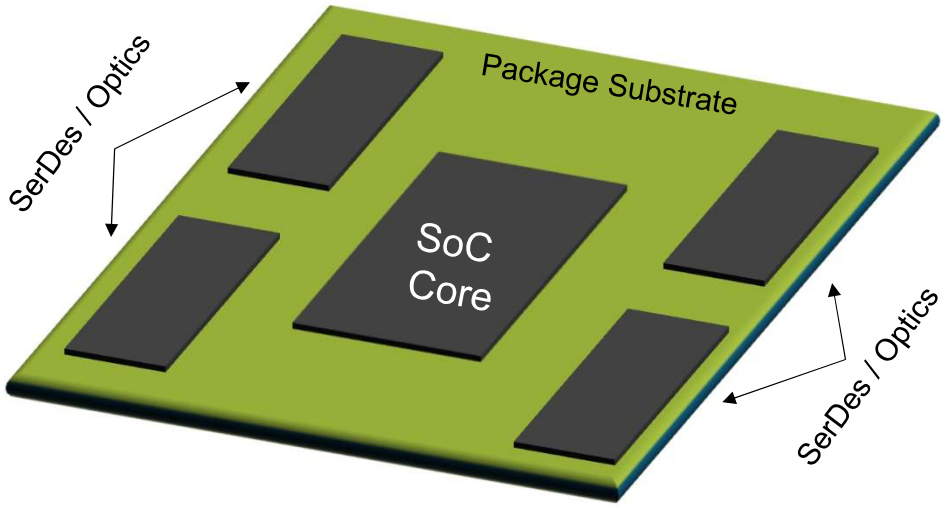

- Typical SiP applications
  - ✓ CPU to CPU in a multi-core SoC, low-latency coherent interconnect
  - ✓ DSP arrays, typically to process information from LIDAR
  - ✓ Switch fabric integration on MCM
  - ✓ Network ASIC to SerDes PMD on separate die
  - ✓ Chip to in-package optical engine

Needs:
- High bandwidth
- Low power
- Low latency
- In-package interconnect

© 2018 Cadence Design Systems, Inc. All rights reserved.

112G/56G Landscape – Datacenter, 5G-Mobile, Optics

Optics

- 50G / 100G
- VSR / LR

Mobile – 5G

- 25G / 50G / 100G
- MR / LR
- RRH / BBU

Datacenter

- 25G / 50G/ 100G
- MR / LR
- Switch/Fabric/NIC

Every specific segment has its own unique PPA and design requirements—one for ALL solutions is NOT going to work

© 2018 Cadence Design Systems, Inc. All rights reserved.
cadence

Cadence SerDes Advantage

Support
Global support team
250M+ parts in production today
100+ support resources

Time To Market
Accelerating time to market in 7FF node
Proven IP in advanced-process nodes

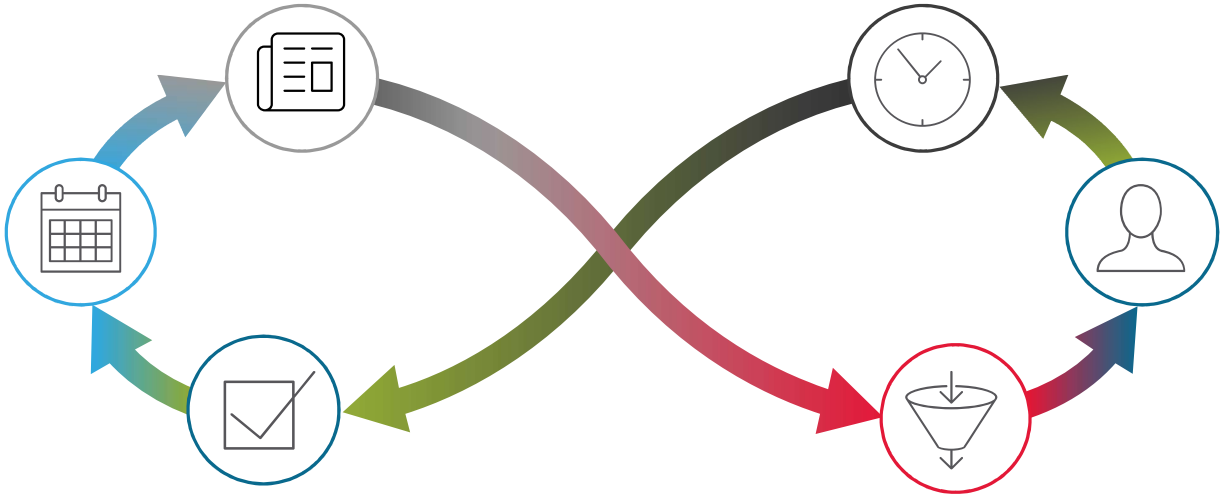

Characterization / Validation
Extensive validation/QA team
Si tests on multiple parts/corners
HTOL, ESD with stress testing

Experienced Design Team
Many years of experience with proven track record of success

Feature Rich Design
Multi-protocol PHY
Fractured clocking architecture
Innovative DFE architecture

PPA
Optimized for protocol efficiency and process nodes

© 2018 Cadence Design Systems, Inc. All rights reserved.

Application-Optimized: Broad Portfolio of PHYs

PPA optimized solution

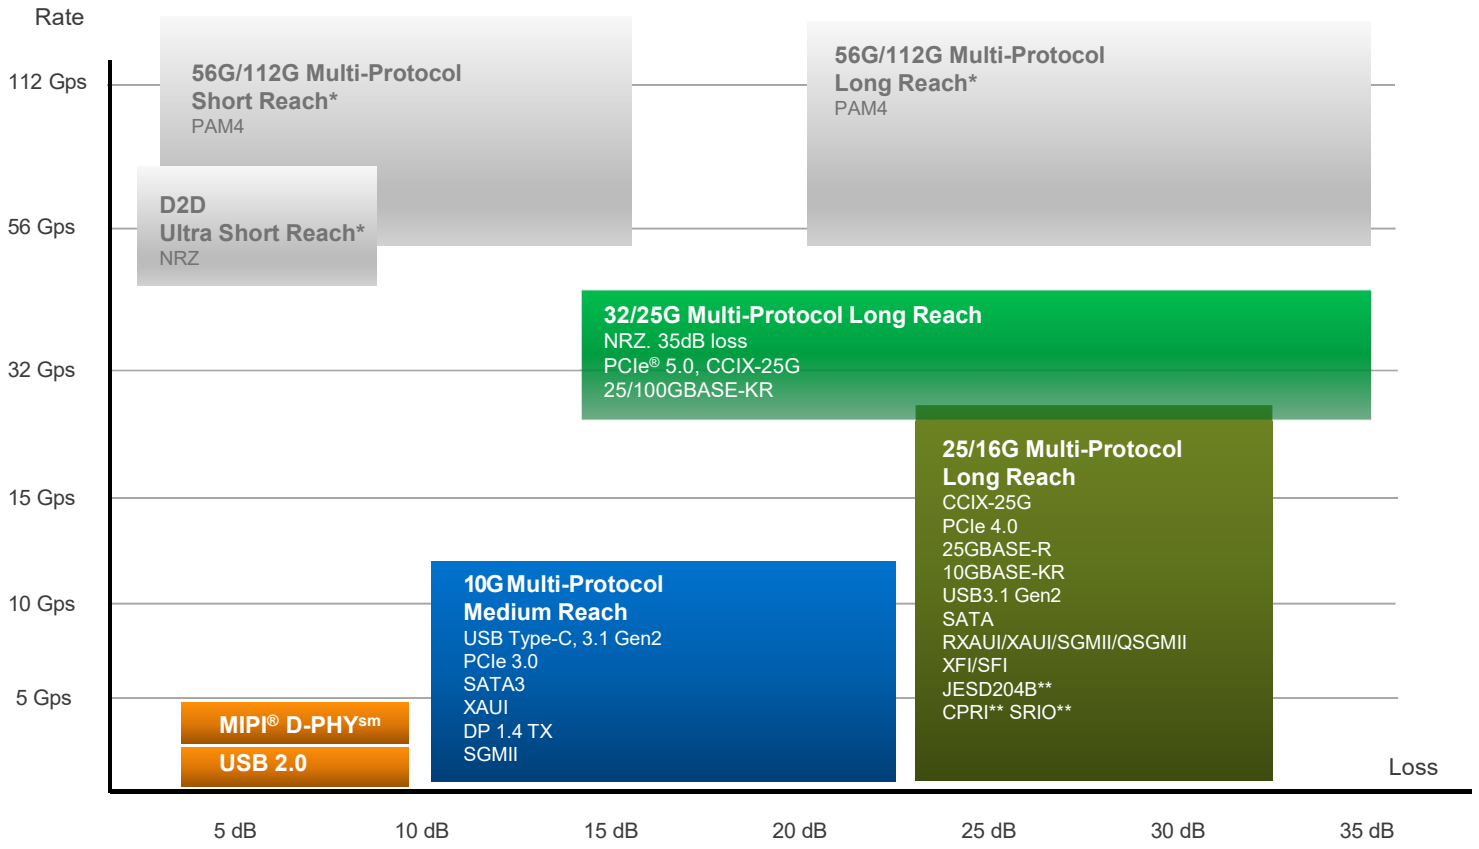

Dedicated PHYs
- USB 2.0
- MIPI-D

Torrent (1-10Gbps)
- Mobile/Consumer
- Low Power
- Small footprint

Sierra (1-25Gbps)
- Mid-End Enterprise
- Multi-Protocol Support
- High Performance
- Low Power

High Sierra (25- 32Gbps)
- High-End Enterprise
- High Performance
- Low Power

* In planning
** Some restrictions apply

This slide contains forward-looking statements regarding *Cadence's business or products. Actual results may differ materially from the information presented here.*

© 2018 Cadence Design Systems, Inc. All rights reserved.

cadence

First to Market: Leader in New Nodes

Advanced technology with proven results

© 2018 Cadence Design Systems, Inc. All rights reserved.

Time to Market: Integrated PCIe Solution

Controller + PHY with drivers and VIP

- Integration of controller and PHY
  - Time-consuming
  - Risky
  - Adds no value to the product
- Advantages of integrated solution
  - Pre-integrated and verified
  - Reduced risk and cost
  - Accelerated TTM
- Custom driver and documentation
  - Customized for each configuration

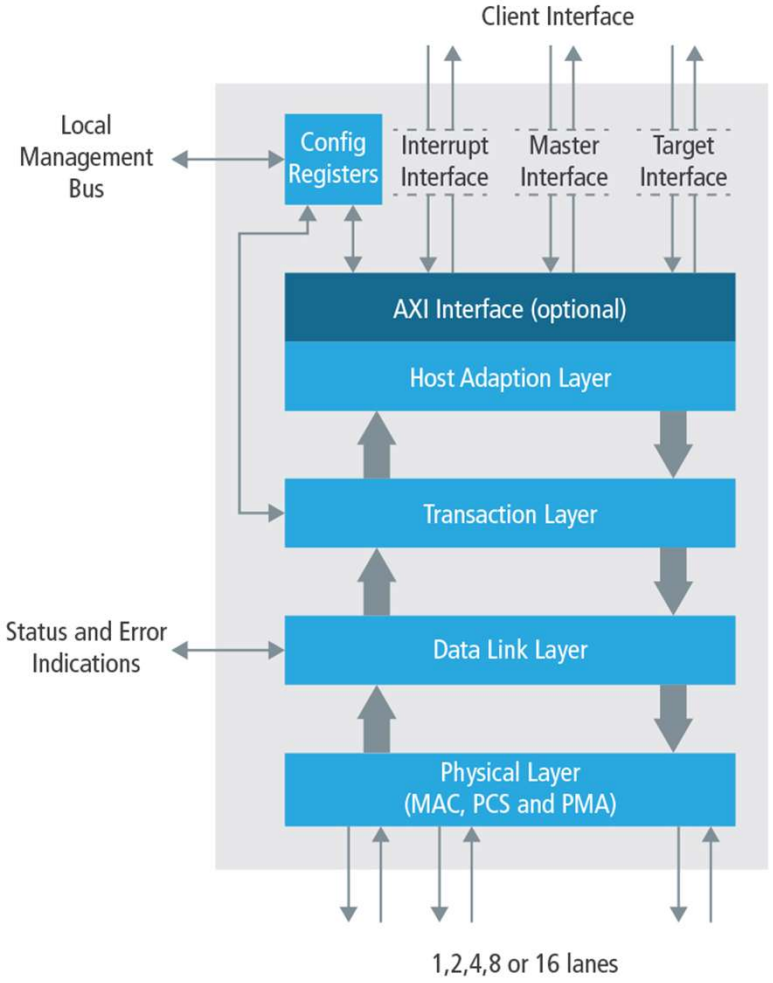

© 2018 Cadence Design Systems, Inc. All rights reserved.

Market Leading Solution: PCIe 4.0 / 5.0 and 32G PHY

Live demo at 2018 PCI-SIG Developer Conference (Santa Clara, USA)

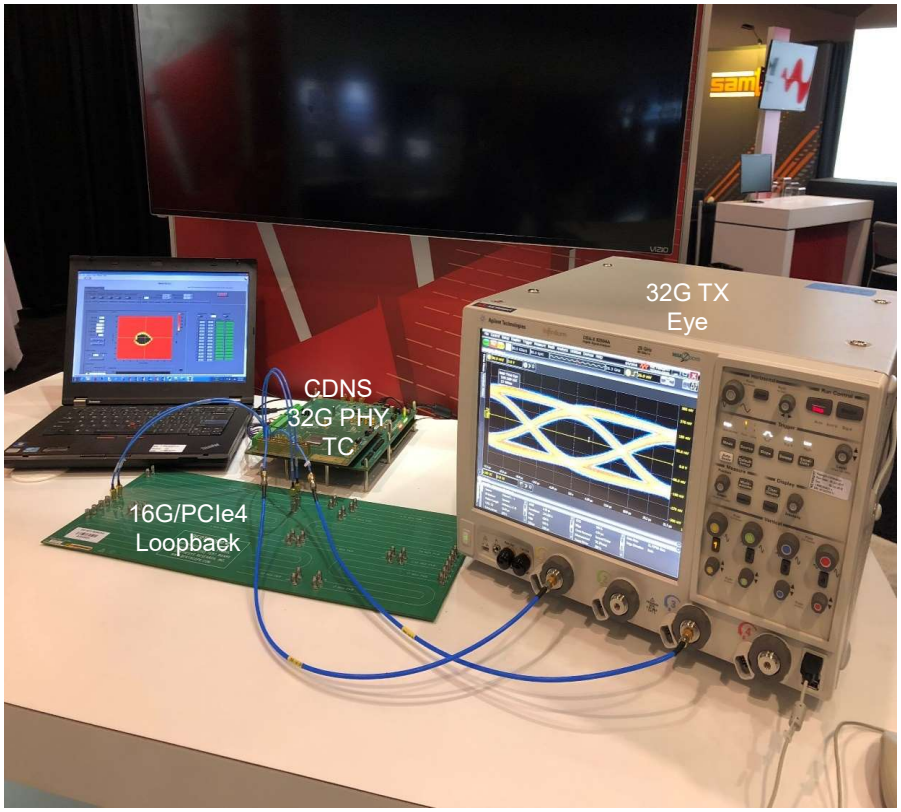

32G Multi-Protocol Multi-Link PHY

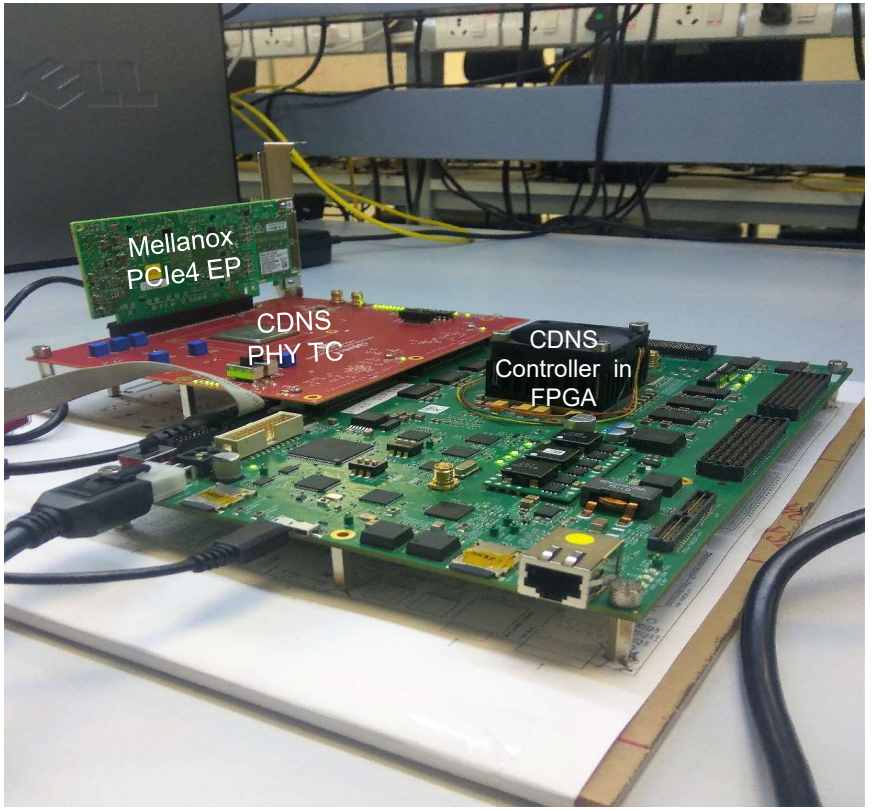

PCIe 4.0 Lane Margining Interop

© 2018 Cadence Design Systems, Inc. All rights reserved.
cadence

Recent Design IP Success

2 of 3 Top Mobile Phone Makers Use Cadence PCIe IP

1st Commercial PCIe Gen4 and CCIX in 7nm

DDR5 7nm Functional in Silicon

17 out of 25 Top Semiconductor Companies Use Cadence IP

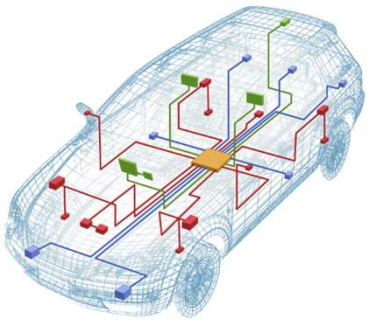

Automobile Leaders Adopting Portfolio in 7nm and 16FFC

© 2018 Cadence Design Systems, Inc. All rights reserved.

cadence

© 2018 Cadence Design Systems, Inc. All rights reserved worldwide. Cadence, the Cadence logo, and the other Cadence marks found at www.cadence.com/go/trademarks are trademarks or registered trademarks of Cadence Design Systems, Inc. Accellera and SystemC are trademarks of Accellera Systems Initiative Inc. All Arm products are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All MIPI specifications are registered trademarks or service marks owned by MIPI Alliance. All PCI-SIG specifications are registered trademarks or trademarks of PCI-SIG. All other trademarks are the property of their respective owners.