Tech Insights

Linley Fall Processor Conference 2022

November 1-2, 2022

Premier Sponsor

flexlogix Technologies, Inc.

Platinum Sponsor

ceremorphic®

Gold Sponsors

CEVA
cadence
ARTERIS IP
MOVELLUS
TetraMem
MARVELL
GlobalFoundries
SYNOPSYS® Silicon to Software
Achronix Data Acceleration
proteanTecs
ANDES TECHNOLOGY
expedera
kinara
UNTETHER AI
quadric
innatera
brainchip
Imagination
SIEMENS
BLUE CHEETAH
Qualcomm
NXP
SiFive
intel

TechInsights

Session 3: Application-Specific Processors

- **Tailored Automotive Radar Processing with DSPs**
  - David Bell, Product Marketing Manager for Tensilica CPU, DSP and AI Products, **Cadence**
- **5G RAN Baseband Platform IP for Cellular Infrastructure Solutions**
  - Eric Cowden, Sr Field Application Engineer, **CEVA**
- **DPUs: Where Will They Go Next?**
  - Manoj Roge, Senior Director, **Marvell Technology**

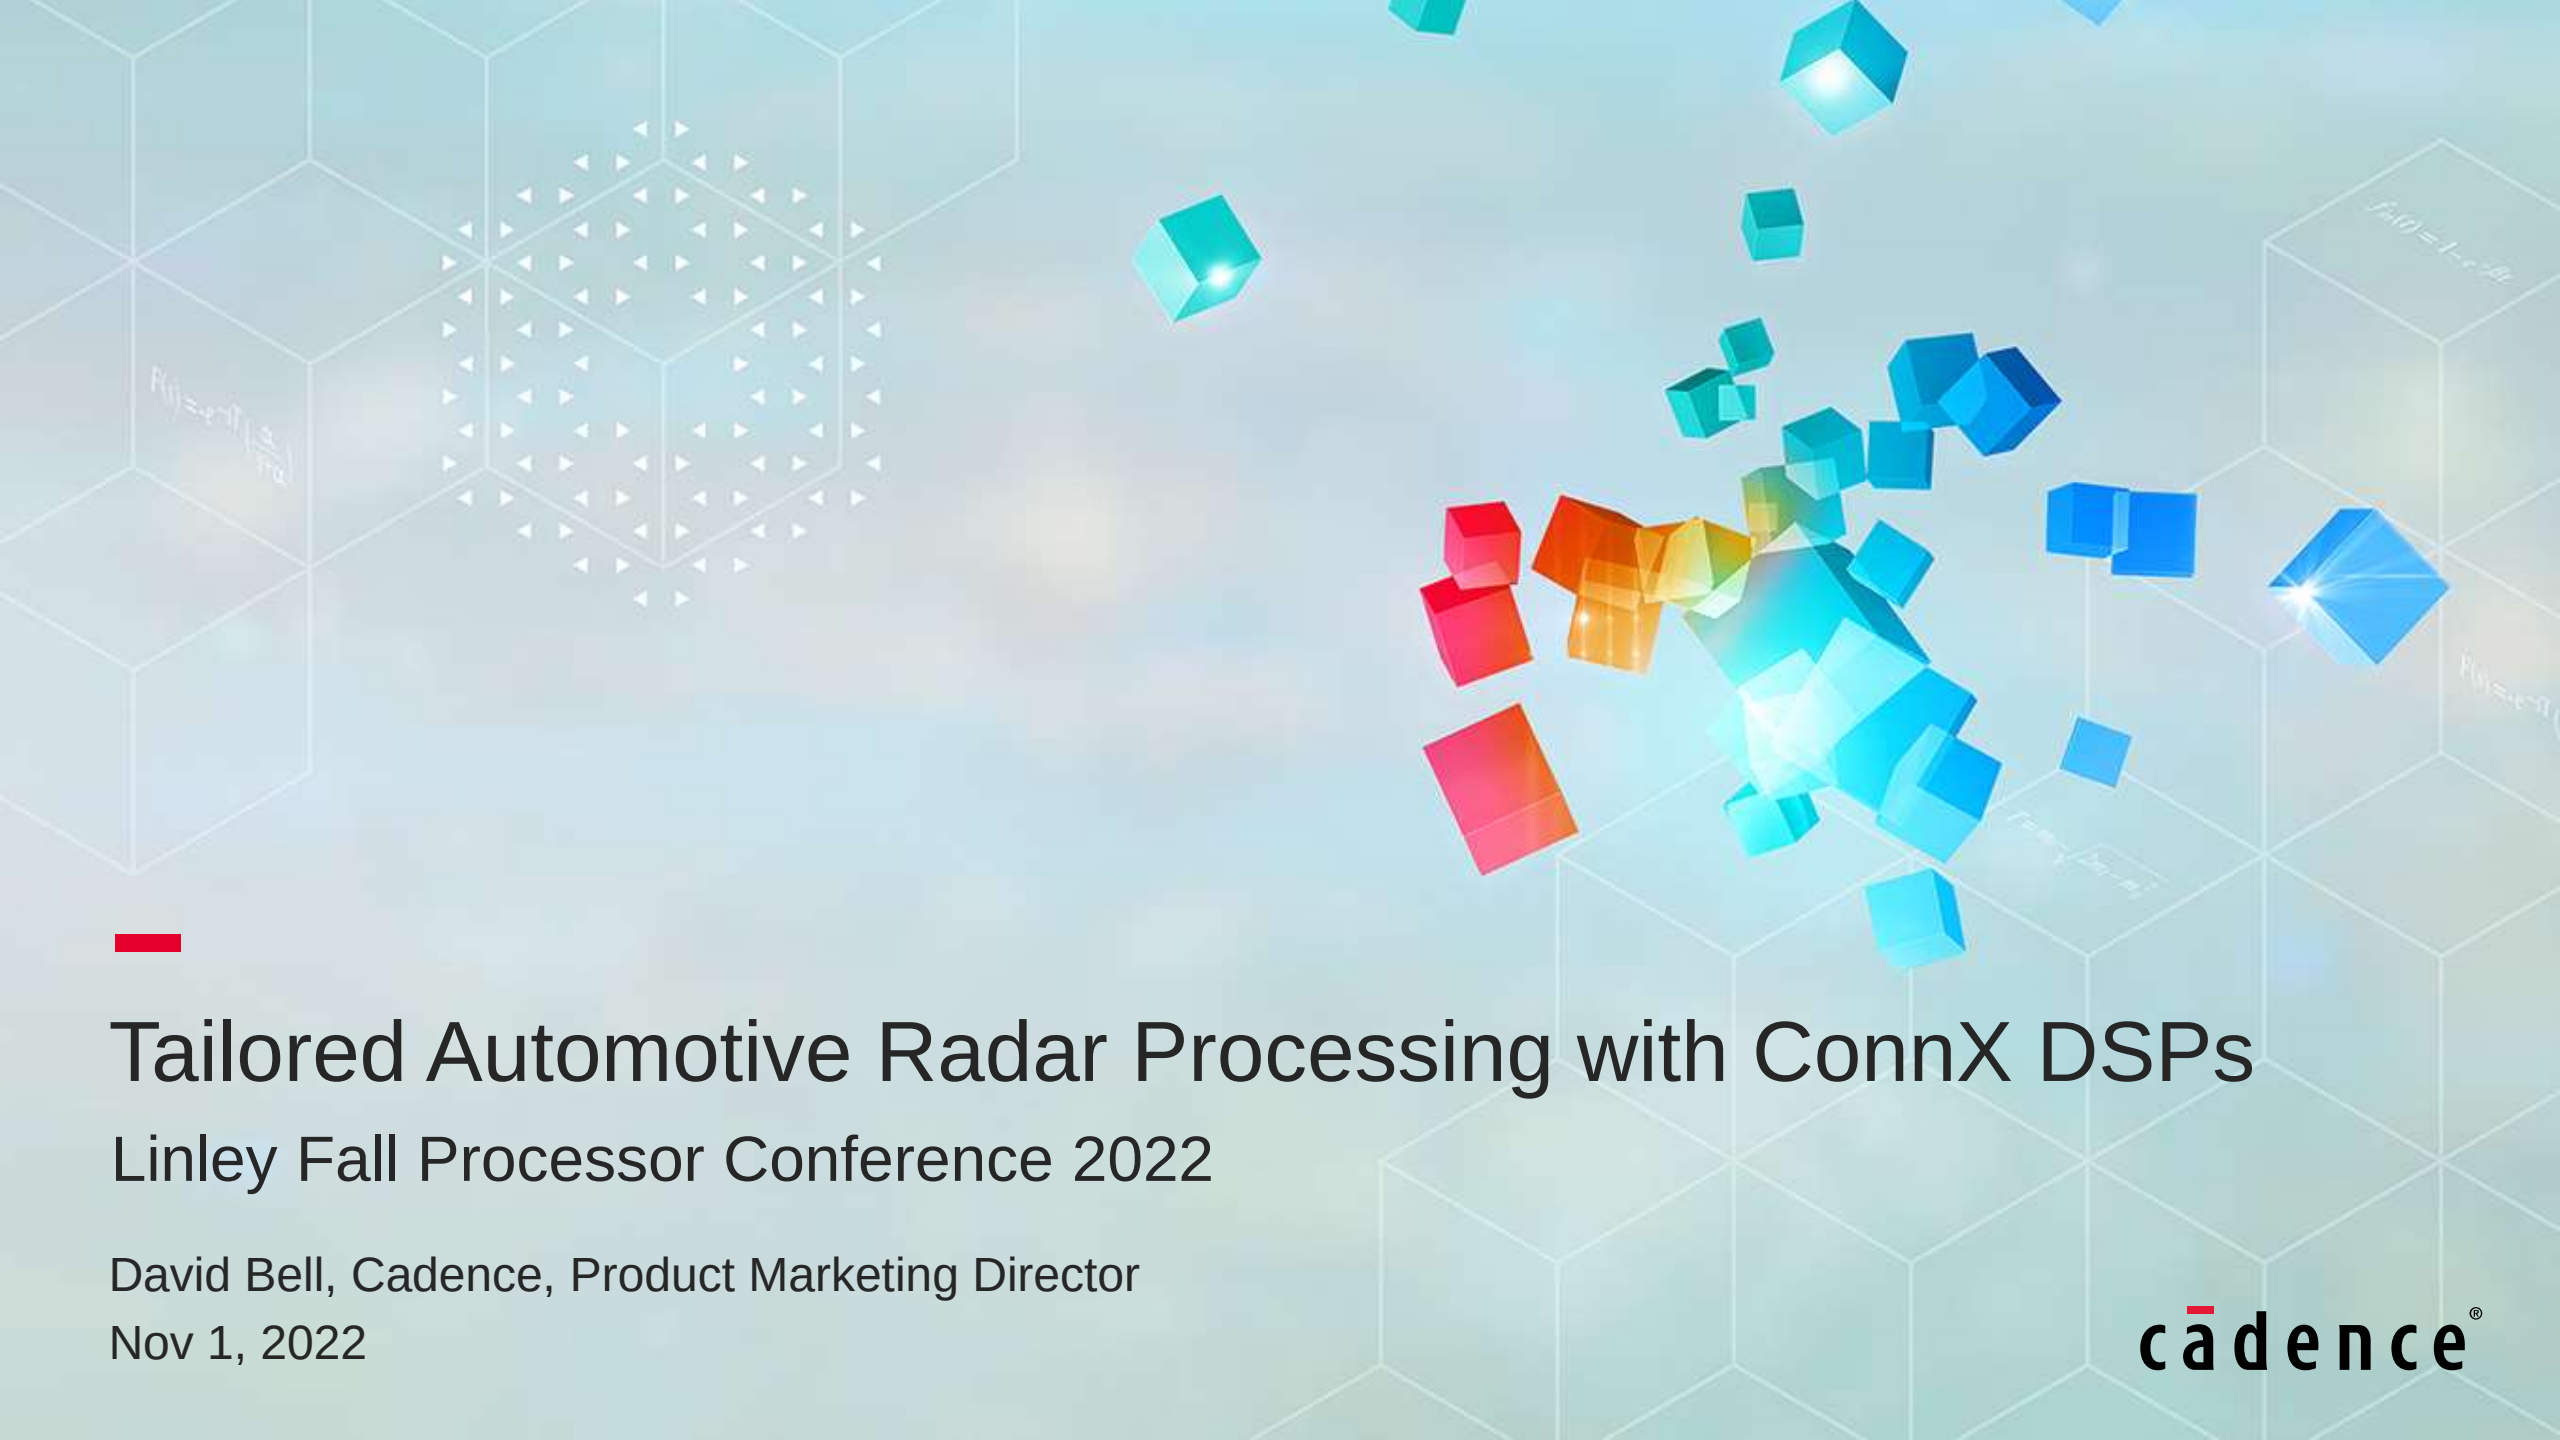

Tailored Automotive Radar Processing with ConnX DSPs

Linley Fall Processor Conference 2022

David Bell, Cadence, Product Marketing Director

Nov 1, 2022

cadence®

Agenda

- Sensors in Automotive
- Radar Processing Requirements
- Tensilica ConnX DSP Introduction
- FMCW Radar Processing Performance with ConnX DSP
- Summary

ADAS Computation Types

Fusion

Radar, LIDAR, Image correlation
System Functional Safety
System Data Control

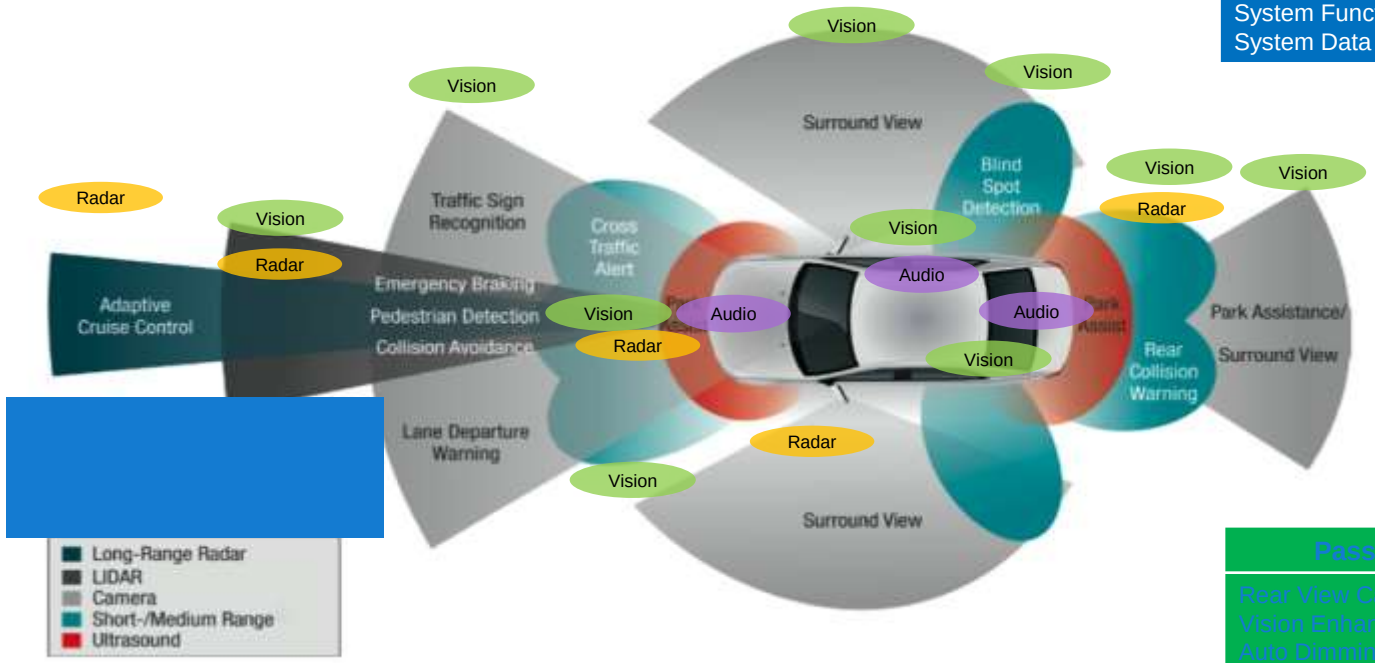

Audio

Rear Object Detection
Parking Assist/Auto Park
Voice Recognition
Cabin Noise Reduction
Emergency Recognition
Spatial Audio for Warnings

Radar

Front Collision Avoidance Braking
Adaptive Cruise Control
360 degree Hazard Awareness
Rear Collision Detection

Active Vision (LiDAR)

Adaptive Cruise Control
Collision Avoidance
Blind Spot Detection

Passive Vision

Rear View Camera
Vision Enhancement
Auto Dimming Headlights
Blind Spot Detection
360 View
Parking Assist
Lane Detection and Following
Sign Recognition
Traffic Signal Recognition
Rain, Snow, /Fog Removal
Pedestrian Tracking /Avoidance
Eye Focus Detection
Driver Monitoring
Vehicle Detection/Avoidance

Source: Cadence

Comparative View of Sensors

Sensor Fusion Key to Autonomous

WCP

No sensor type works well for all tasks and in all conditions, so sensor fusion will be necessary to provide redundancy for autonomous functions

Most likely used fusion solution in future — Good — Fair — Poor

	Camera	Radar	LiDAR	Ultrasonic	LiDAR+Radar+ Camera
Object detection	Fair	Good	Good	Good	Good
Object classification	Good	Poor	Fair	Poor	Good
Distance estimation	Fair	Good	Good	Good	Good
Object edge precision	Good	Poor	Good	Good	Good
Lane tracking	Good	Poor	Poor	Poor	Good
Range of visibility	Fair	Good	Fair	Poor	Good
Functionality in bad weather	Poor	Good	Fair	Good	Good
Functionality in poor lighting	Fair	Good	Good	Good	Good

"Sensor fusion is key because the more complex features get, the more redundancy you need. Every autonomous vehicle is going to have some combination of LiDAR, Radar and camera."
– ADAS engineer at a prominent OEM

Source: Expert interviews

WOODSIDE CAPITAL PARTNERS

Lidar

Radar

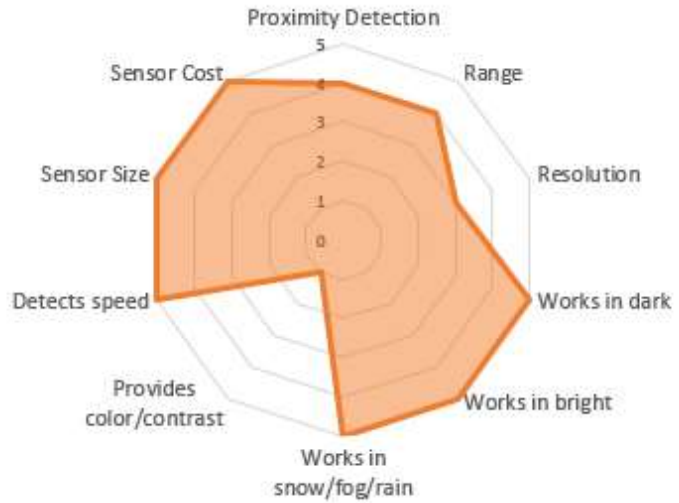

cadence

Automotive Sensors Market

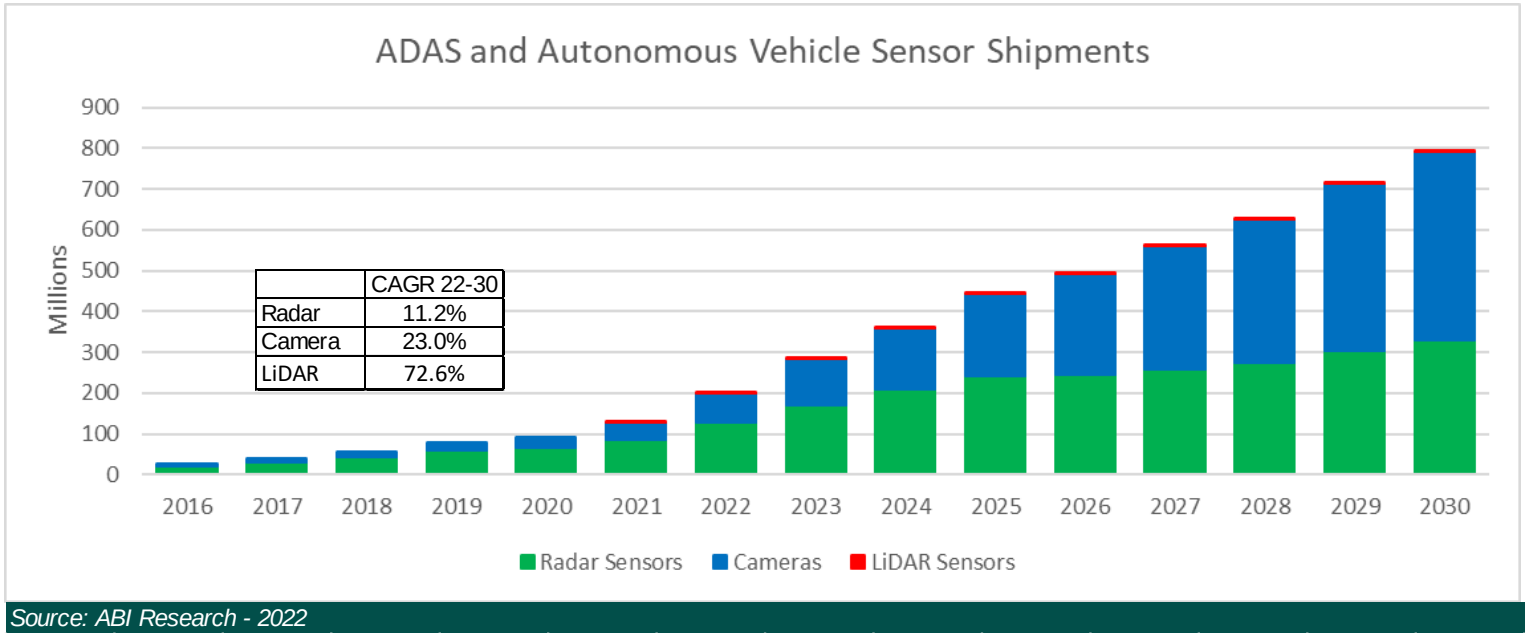

Source: ABI Research - 2022

Radars	Uses
Long Range	ACC
Short/Medium Range	Emergency Braking, Cross traffic alert, Blind spot detection, Rear collision warning, occupancy sensing, etc.

System	Sensors
Forward Collision Warning	Radar, Camera, LiDAR
Forward Collision Warning and Mitigation/Active Emergency Braking	Radar, Camera, LiDAR
Adaptive Cruise Control	Radar, Camera
Traffic Jam Assist	Radar, Camera
Intelligent Speed Assist	Camera
Blind Spot Detection	Radar, Camera
Lane Change Assist	Radar
Lane Departure Warning	Camera
Lane Keep Assist	Camera
Parking Assist	Ultrasonic Sensors
Active Parking	Ultrasonic Sensors, Camera
360° View/Reversing Camera	Camera
Cross Traffic Assist	Radar, Ultrasonic
Pedestrian Detection	Camera, Radar, Infrared
Cyclist Detection	Camera, Radar, Infrared
Motorcyclist Detection	Camera, Radar, Infrared
Animal Detection	Camera, Radar, Infrared
Attention Assist	Camera, Steering Wheel, Other Vehicle Sensors, Biometric Sensors
Fatigue Detection	Camera, Steering Wheel, Other Vehicle Sensors, Biometric Sensors
Night Vision	Infrared
Adaptive Headlights	Camera, Steering Wheel
High Beam Assist	Camera
Road Sign Recognition	Camera
Driver Monitoring Systems	Camera

Purpose of a Radar and Classification

- Range determination: How far is the target?
- Velocity determination: How fast is the target moving relative to the radar?
- Angle determination: In which direction is the target moving?

- Some radars determine some of the above features
  - Frequency Modulated Continuous Wave (FMCW) Radar can determine all three
    - Maximum range determinable, maximum velocity determinable, maximum Field-of-View determinable
    - Range resolution, velocity resolution and angle resolution

- Classification of automotive radar based on system properties as
  - Short-, Medium-, and Long-Range Radars

Basic Principle of a Radar

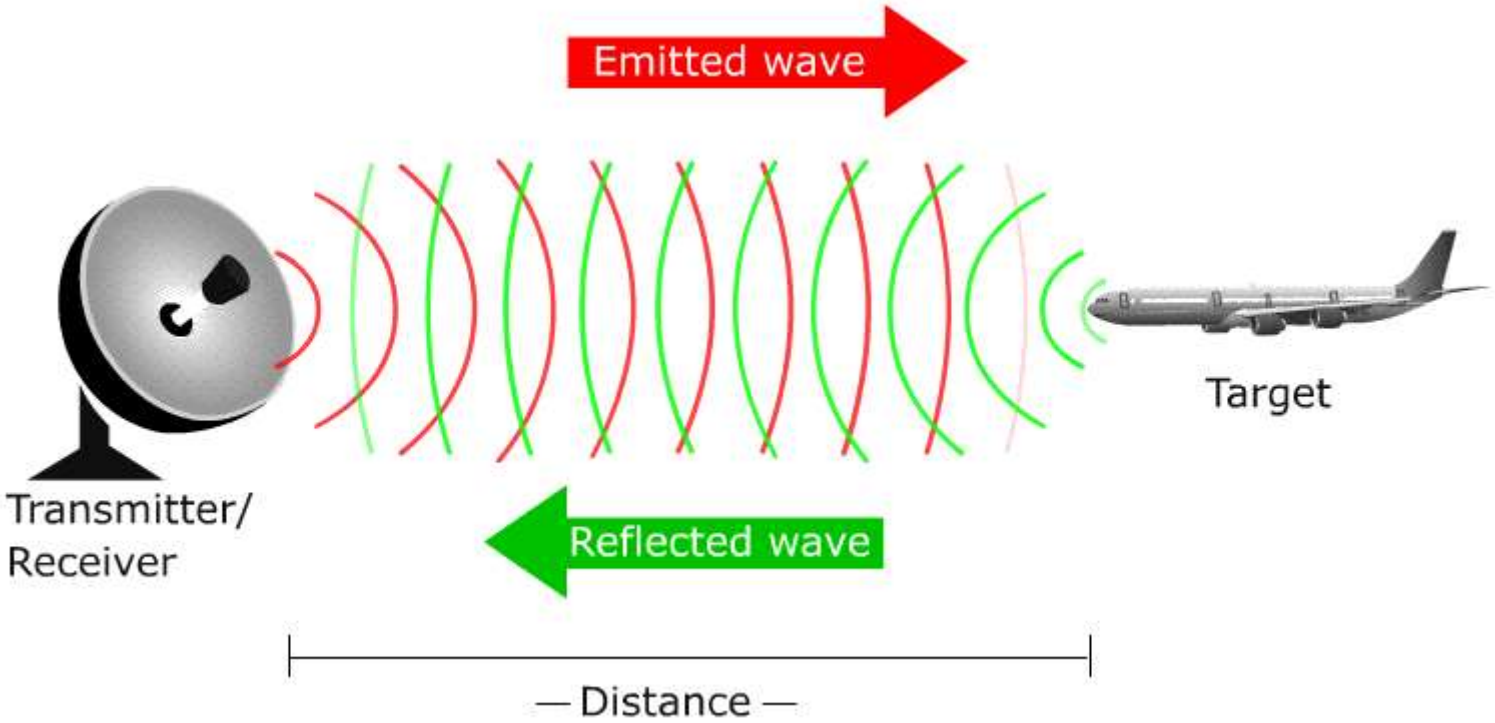

Range determination: FMCW Chirp and Intermediate Frequency (IF)

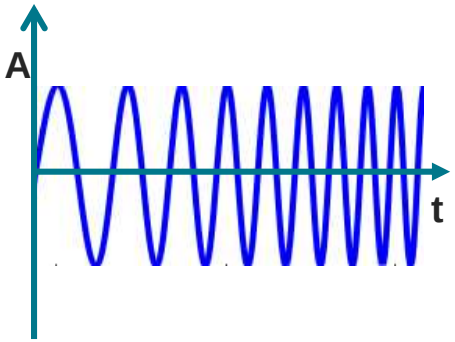

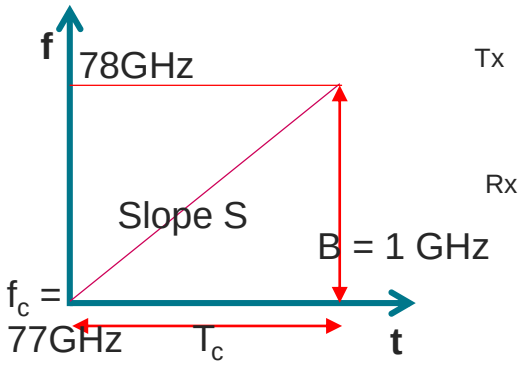

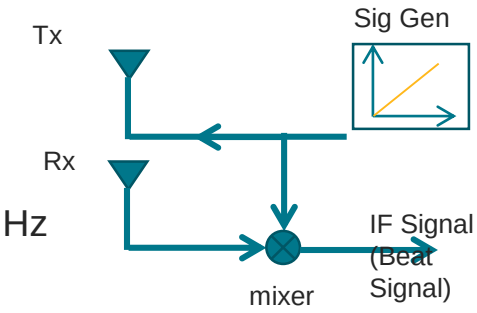

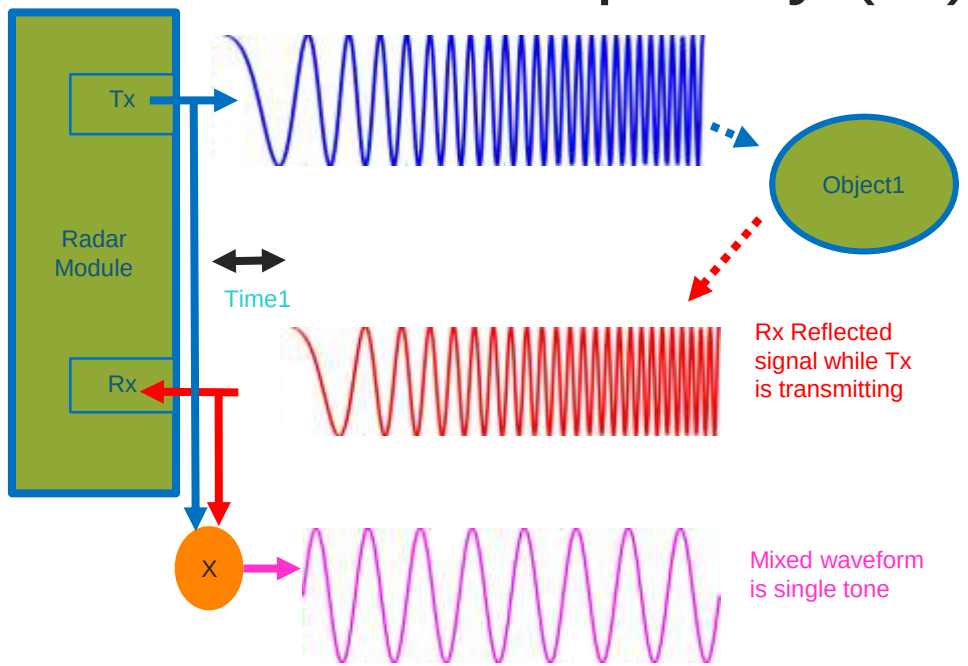

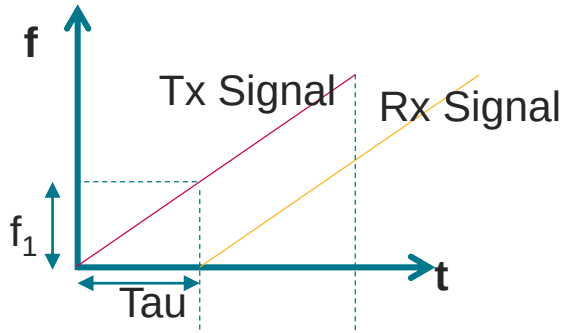

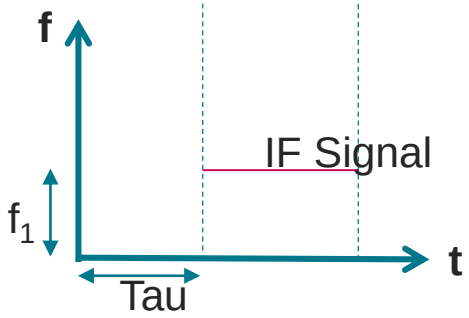

$$f_1 = S * Tau = S * 2 * d/c$$

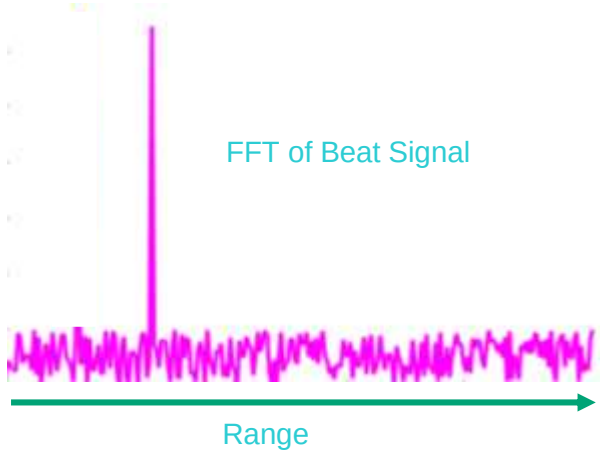

f_c	Carrier Freq
T_c	Time period of chirp
B	Bandwidth of chirp
S	Slope of chirp
Tau	Round trip delay
d	Distance of target
c	Speed of light
f_1	IF corresp. to target 1

cadence

Velocity Determination: How to Measure Velocity

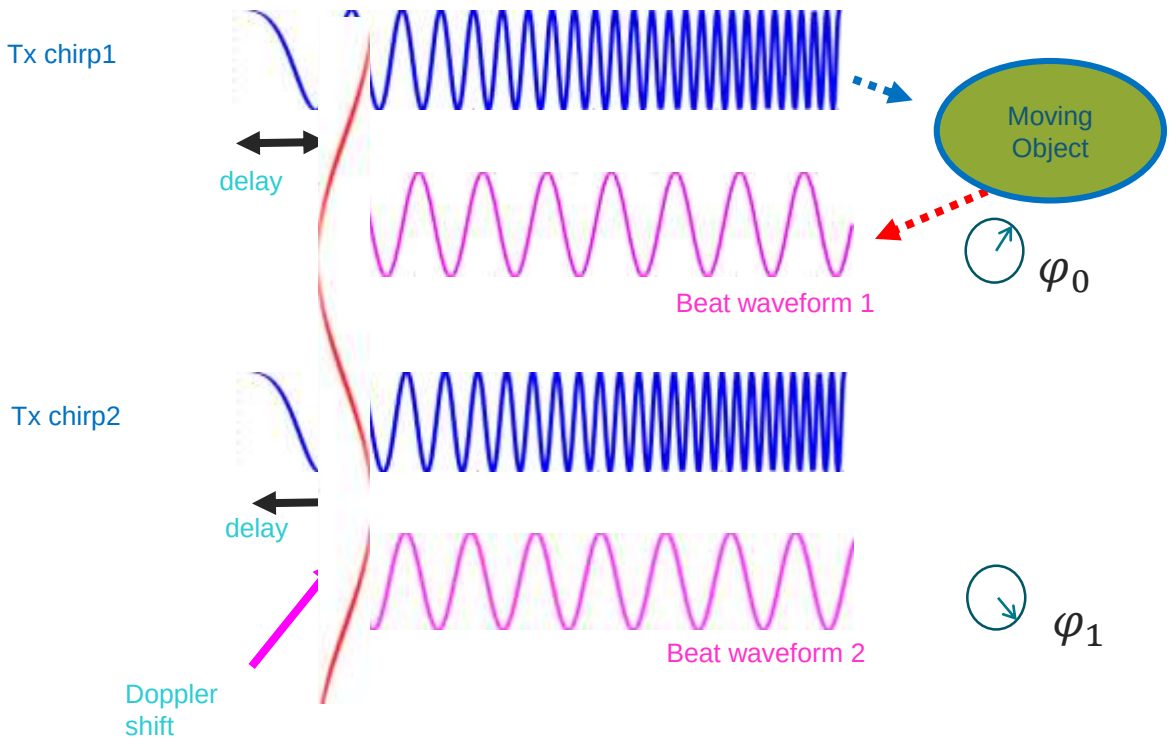

- To measure the velocity of a target, transmit two chirps
- The change in phase of the IF signal across the chirps would correspond to the velocity (change in target location in a single chirp duration) of the target
- Phase difference between chirps

$$\omega = 4\pi v T_c/\lambda$$
$$v = \lambda\omega/4\pi T_c$$

Angle: Angle Determination

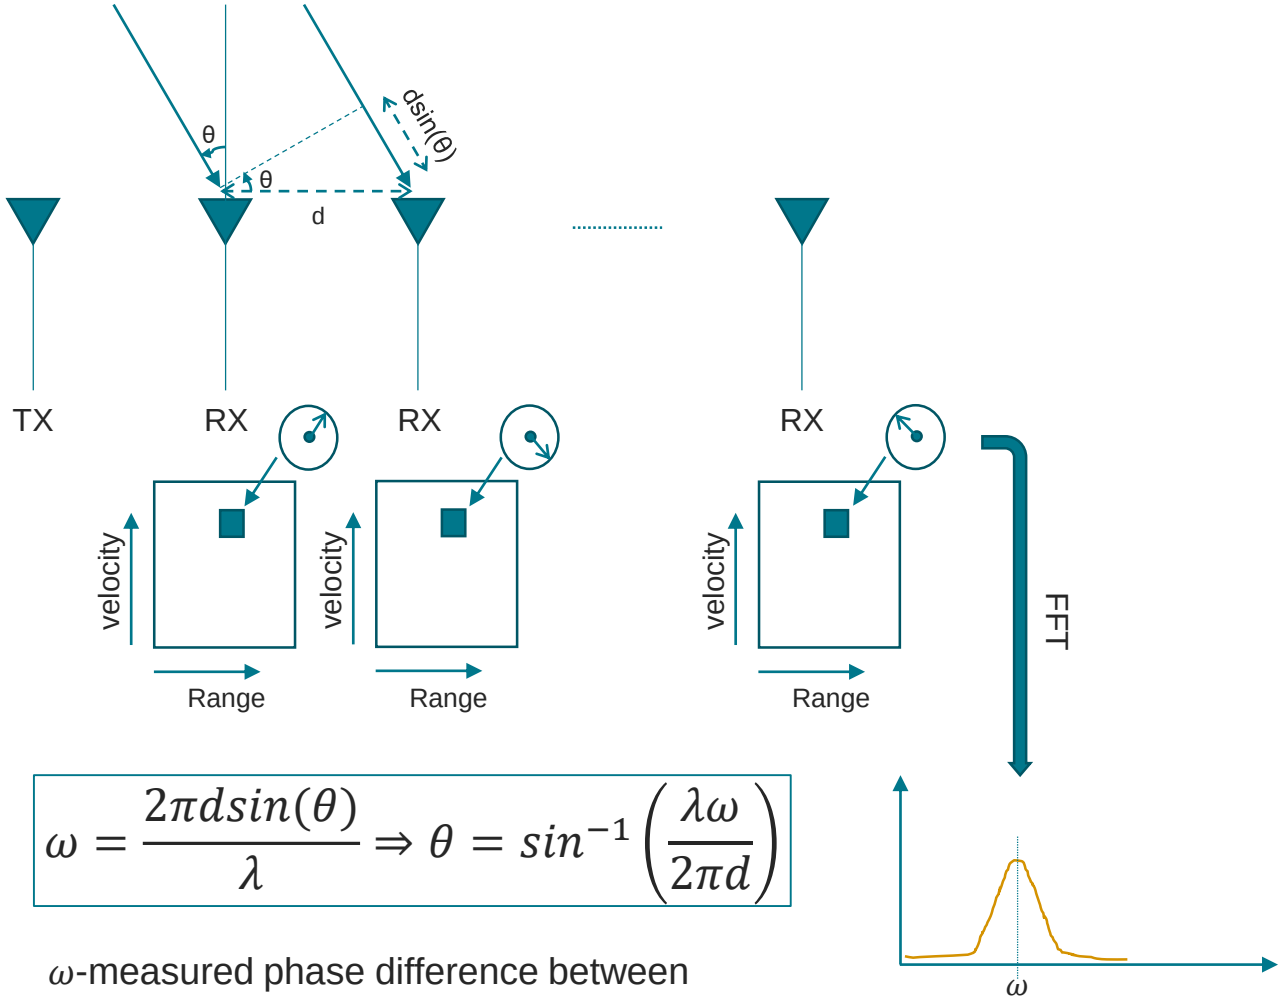

$$\omega = \frac{2\pi d sin(\theta)}{\lambda} \Rightarrow \theta = sin^{-1}\left(\frac{\lambda\omega}{2\pi d}\right)$$

ω-measured phase difference between two antennas

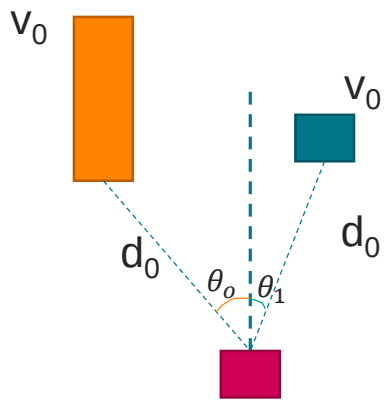

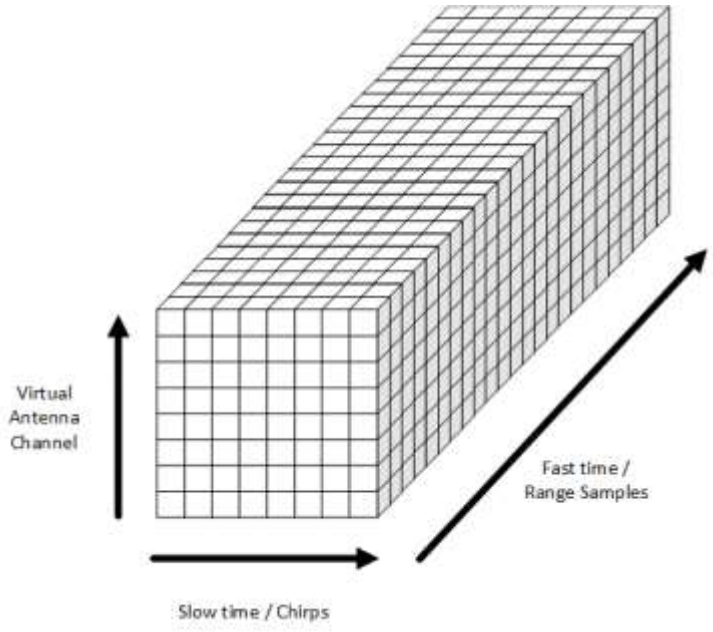

Radar data cube

FMCW Radar Signal Processing Chain

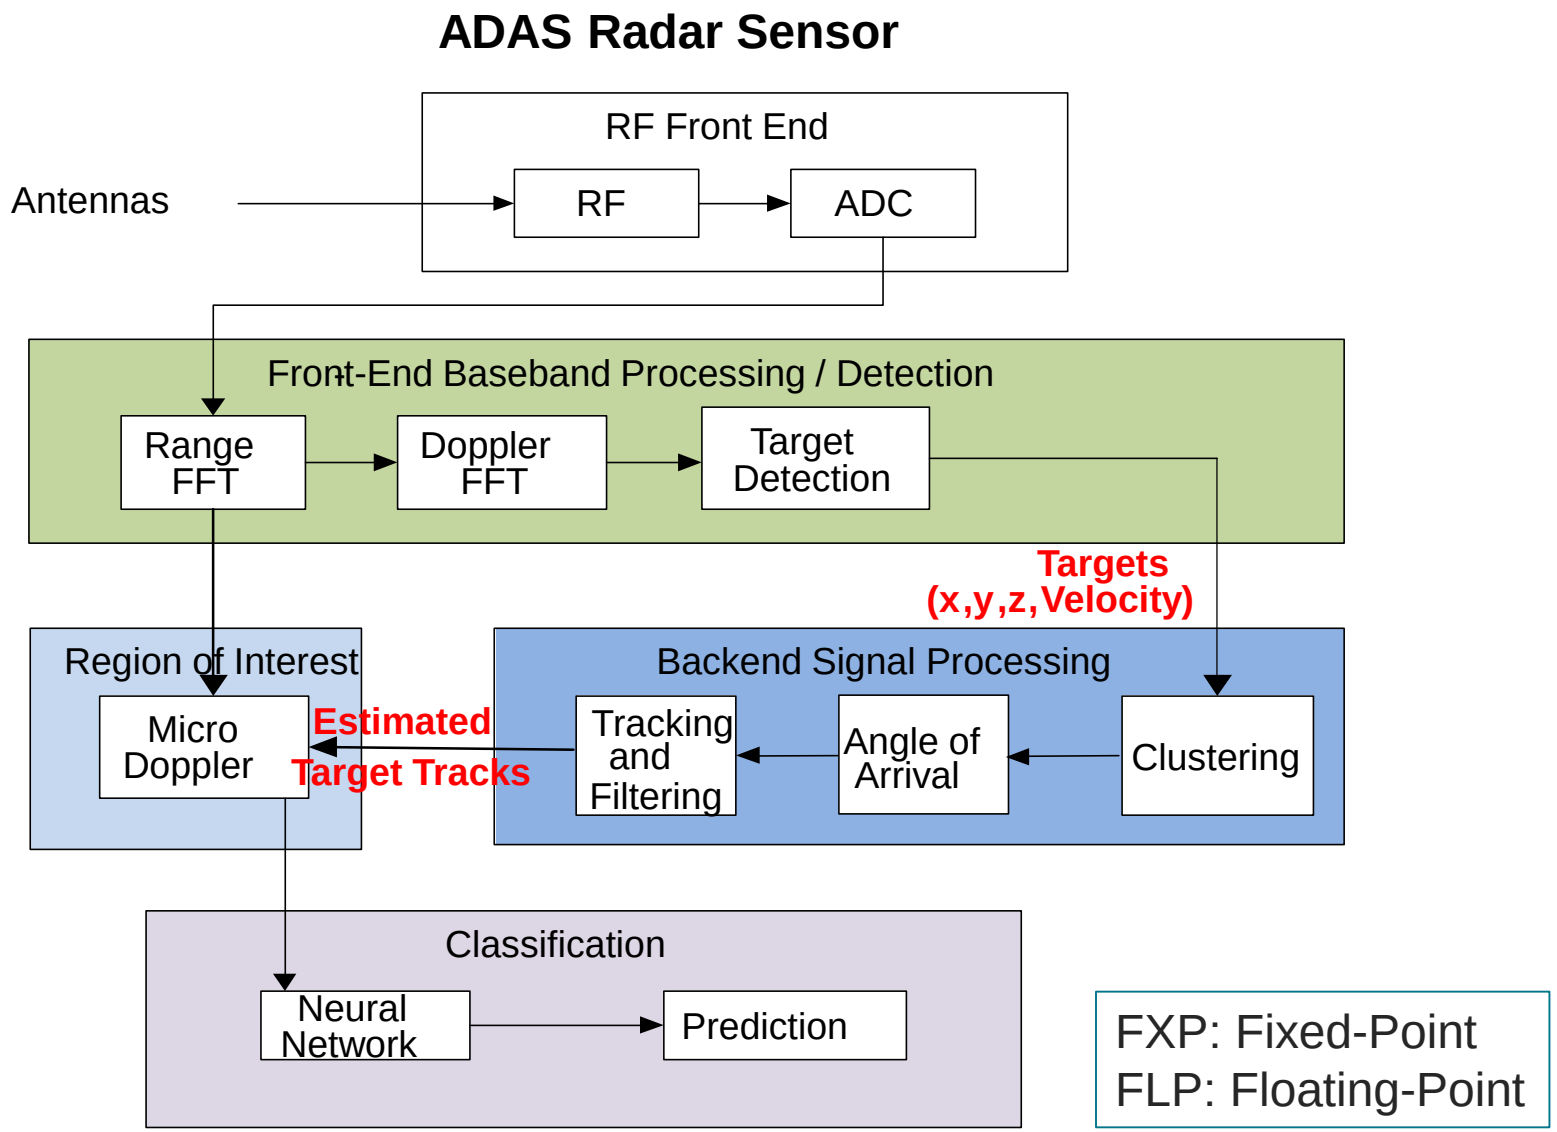

FXP: Fixed-Point
FLP: Floating-Point

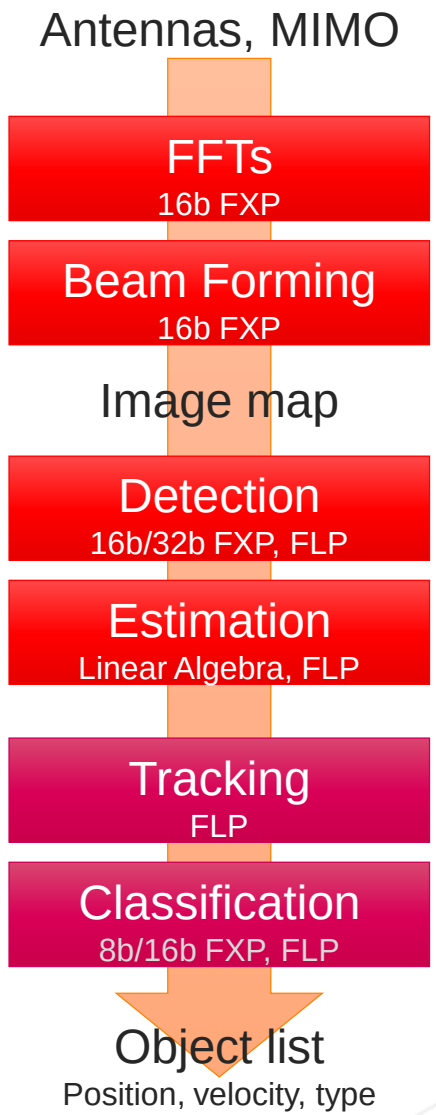

Tensilica ConnX DSPs

Automotive, Consumer, Industrial

- ConnX B20
  512-bit SIMD
- ConnX B10
  256-bit SIMD
- ConnX 120
  256-bit SIMD
  New
- ConnX 110
  128-bit SIMD
  New

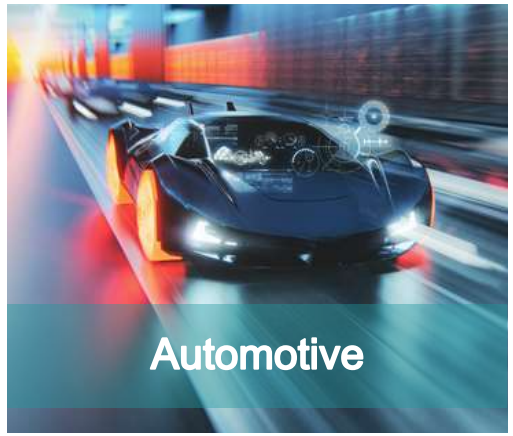

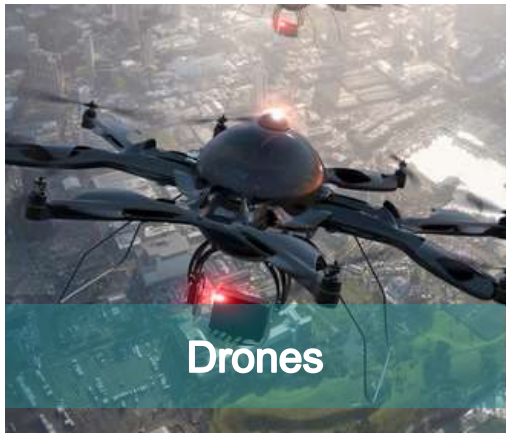

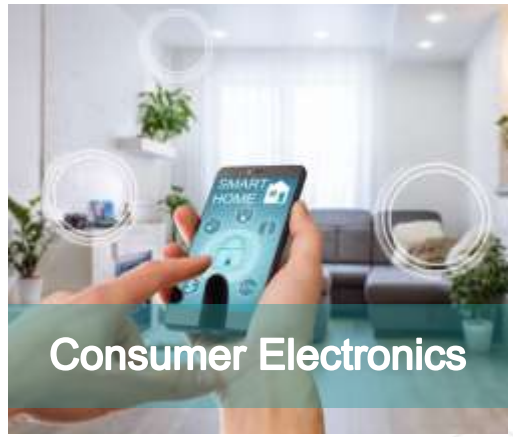

cadence

Tensilica ConnX DSP Family

Scalable, configurable, extensible, flexible

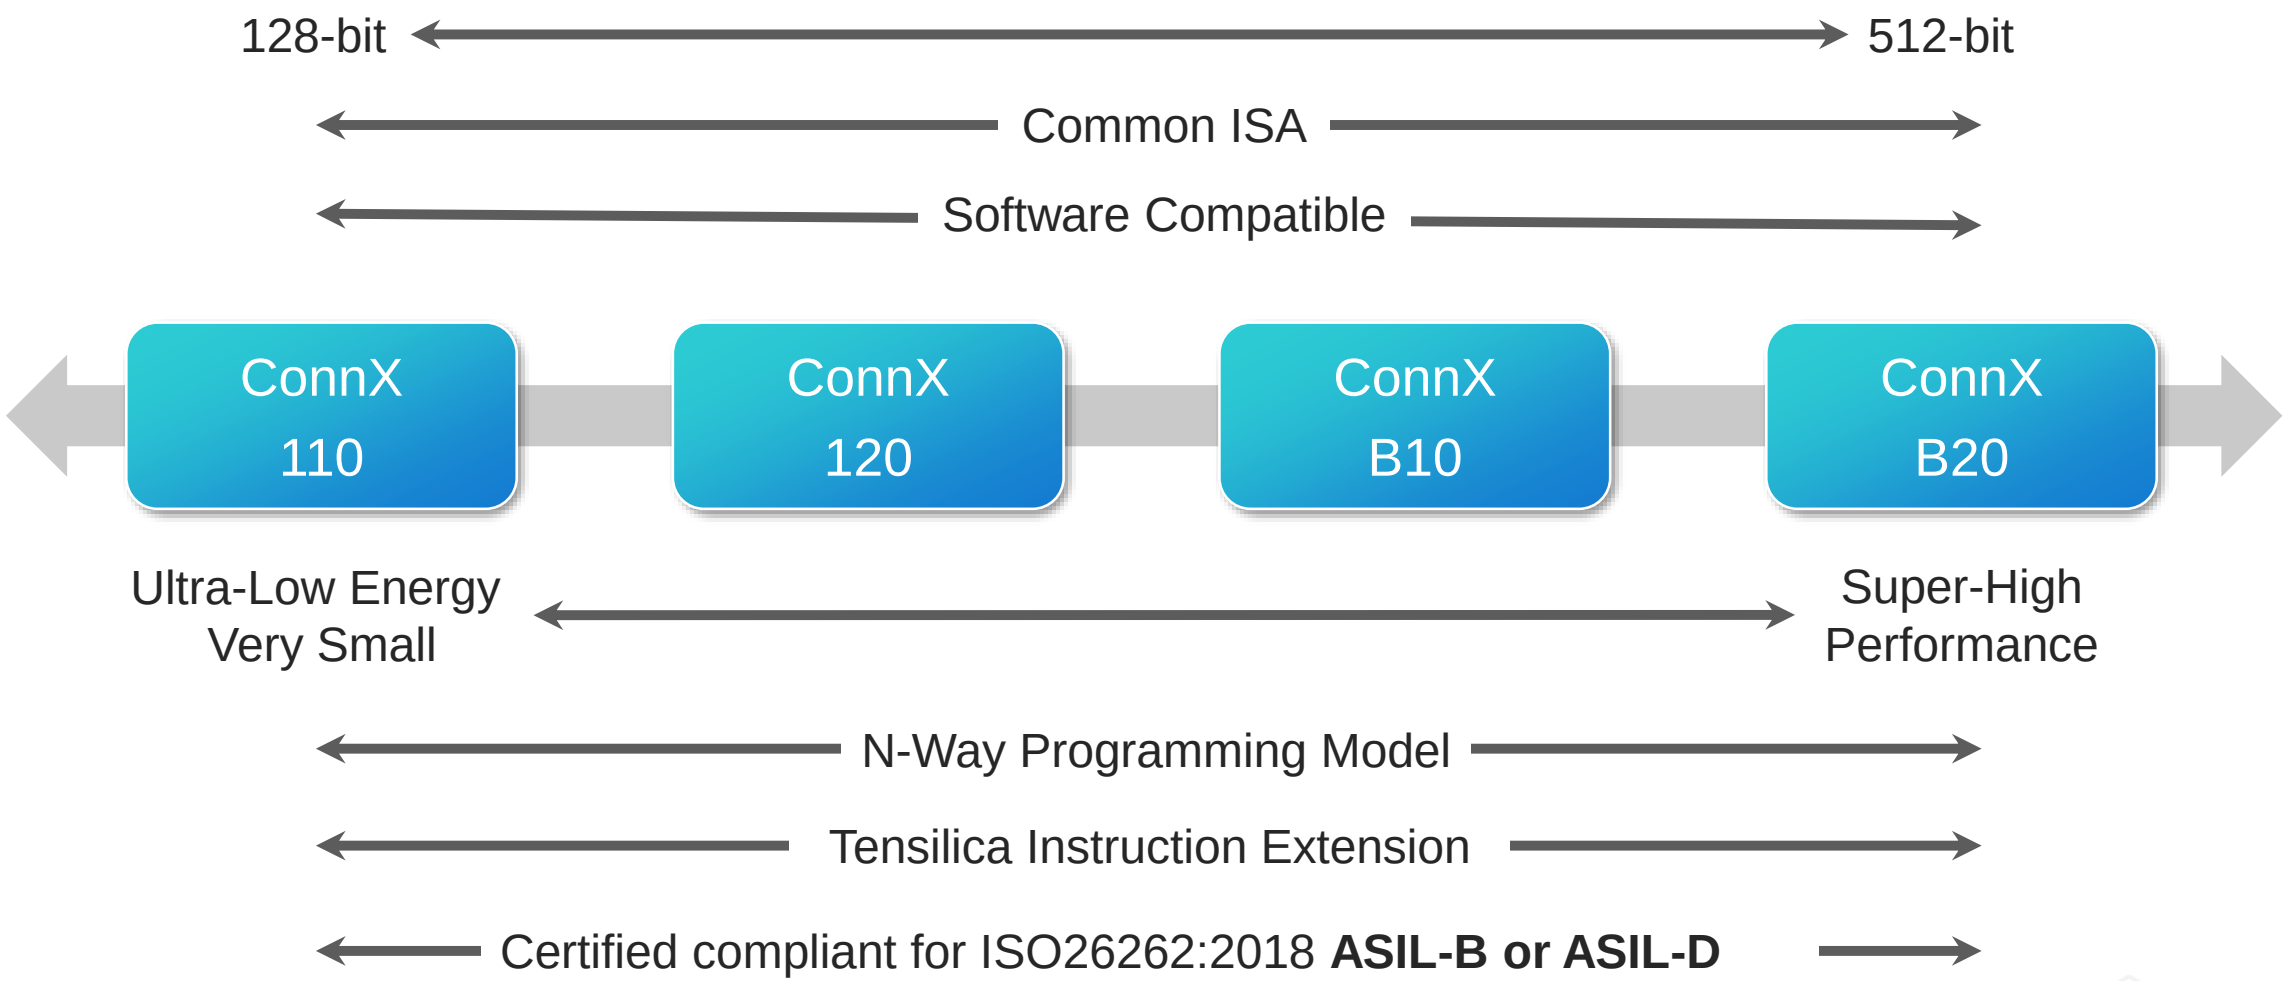

Radar – Customer Success

- Renesas RH850/V1R-M
  - Cadence® Tensilica® ConnX BBE32 DSP
  - ISO26262: Functional Safety

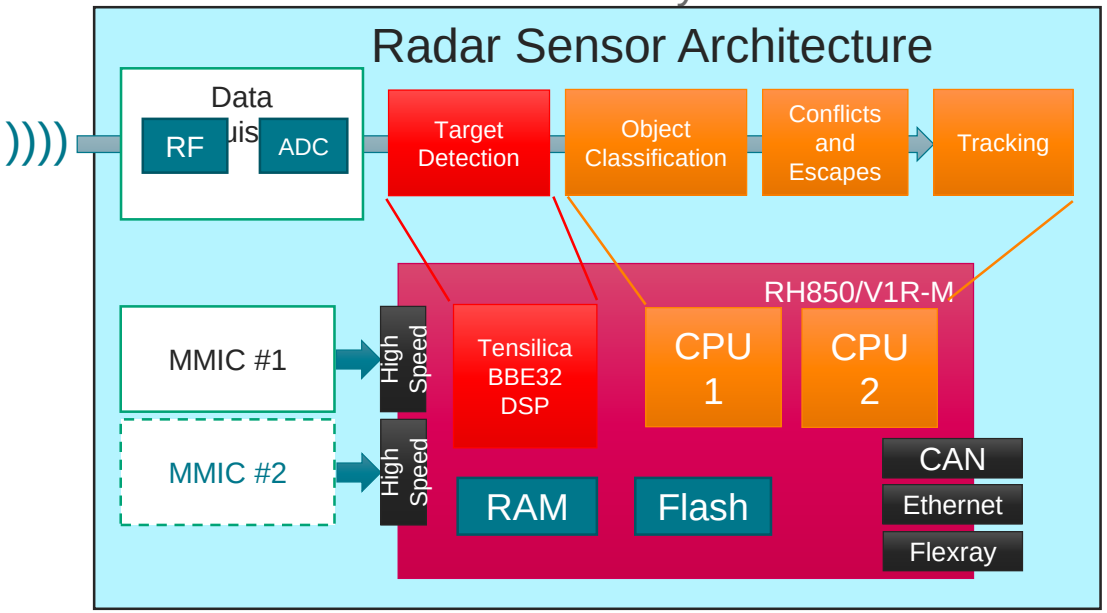

  - References
    - Renesas: Balancing performance, low-power and functional safety in ADAS Applications (7/8/2016)
    - Renesas Radar Solution RH850/V1R-M (12/23/16)

- NXP
  - ConnX BBE32EP DSP-based solution with custom instructions (TIE)
  - "*…Specifically in radar, the contribution of Cadence has been fundamental. The [enhanced] DSP capability [allowed] us to develop algorithms that see with more accuracy, farther and to avoid ghosts or false positives...*"

Source: https://youtu.be/LEwUfTs82lI

ConnX High-Performance DSPs

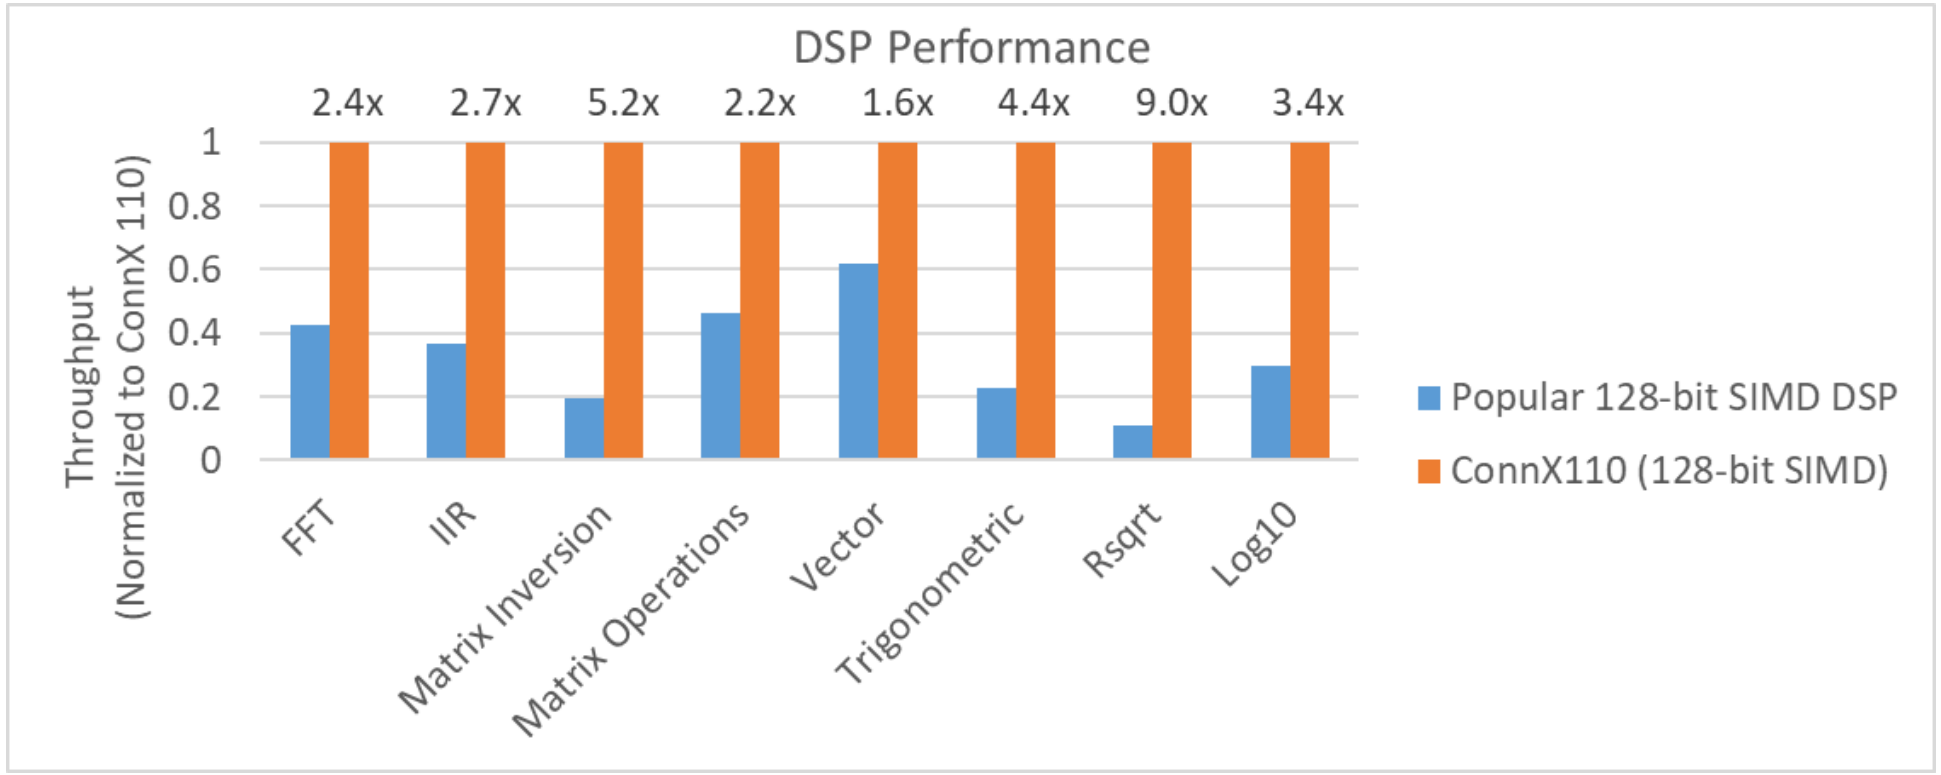

- ConnX DSP provides high-performance processing across DSP and Math functions essential to radar processing
  - ConnX 120 and B10 is ~2X the performance of ConnX 110
  - ConnX B20 is ~4X the performance of ConnX 110

Radar Trend Implications

- Detection requirements are increasing
  - Number of objects, relative speeds, max distances and speeds, resolution
- System/algorithm requirements/constraints are increasing
  - Number of channels/sensors, SIMO or MIMO
  - Performance, precision, convergence, flexibility
  - Memory hierarchy and accessibility

Module	Type of computation and trend requirements
DBF Beam forming	Dot products > Faster with wider SIMD
RFFT, DFFT Range-Doppler processing	FFT > Higher-precision complex FFT
Constant False Alarm Rate (CFAR) Detection	Averaging, sorting > Fast Peak Searches, Scatter Gather, Complex support
Multiple Signal Classification (MUSIC) High-Res Direction estimation algorithm	Matrix Eigen decomposition, polynomial root finding > Good Floating Point Linear Algebra
Clustering Tracking	Tracking: Kalman Filter (Matrix multiplies/inversion) Clustering: DBSCAN > Good Linear Algebra, Scatter Gather

Software Libraries and Tools for ConnX DSPs

Optimized DSP functions for radar applications

- Scalar/Vector fixed/floating-point, real/complex support
- Library source code available

Library / Tool	Functions
NatureDSP	Optimized Math functions for Logarithmic, Complex, FFT, FIR, IIR, Vector, Matrix operations
Eigen Library	High-level library for linear algebra, matrix, vector, numerical solvers
Radar Library	Range FFT, Doppler FFT, Constant False Alarm Rate (CFAR), Extended Kalman Filter (EKF)
Matlab Toolkit	Simulink and Matlab code generation support to compile and simulate with ConnX DSP targets directly from Matlab

https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/ip/tensilica-ip/tensilica-dsp-code-generation-toolbox-with-matlab-simulink-wp.pdf

FMCW SIMO Example Performance

ConnX B10/B20 – Fixed/Floating Point

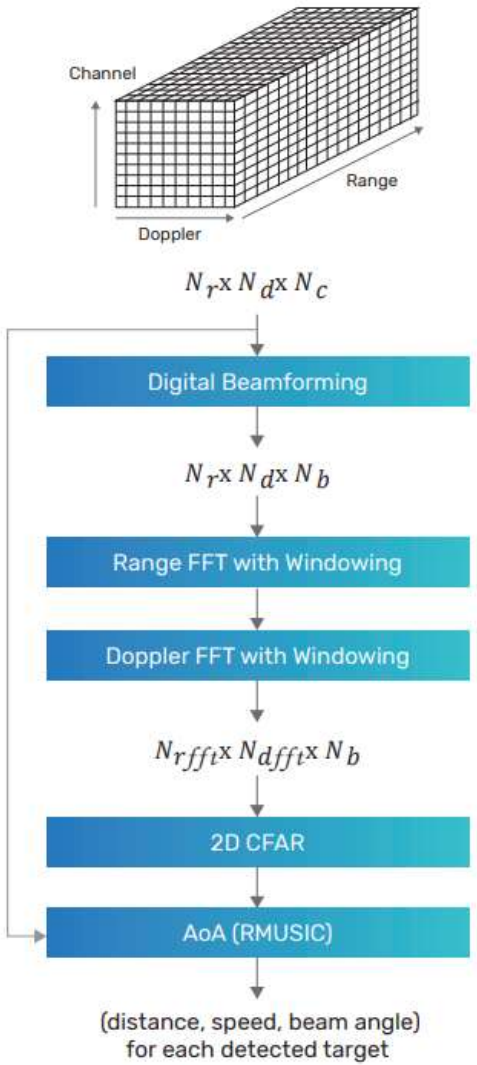

Description	Mixed-Precision	Floating-Point Only
Input Data Cube Size	1024x64x16	
Digital Beamforming (DBF)	Complex 16-bit fixed-point samples (represented in Q15 format)	Complex floating-point samples
Range-Doppler FFT (RFFT and DFFT)	Complex 16-bit fixed-point samples	Complex floating-point samples
Constant False Alarm Rate (CFAR)	Real 32-bit fixed-point	Real floating-point samples
Root Multiple Signal Classification (RMUSIC)	32-bit complex floating point	

Tensilica ConnX DSP	Mixed-Precision	Floating-Point Only
ConnX B10	133MHz	259MHz
ConnX B20	84MHz	154MHz

https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/ip/tensilica-ip/tip-fmcw-radar-wp.pdf

Summary: Choose Tensilica DSPs for Automotive Radar

Fast development	• Optimized software libraries for a wide variety of DSP functions in C • Matlab Toolkit Integration • Select from predefined hardware for algorithm acceleration

Processing flexibility	• Mixed Precision
	  ○ 16- and 32-bit Fixed Point
	  ○ Half-, Single-, Double-Precision Floating Point

Scalability	• Code-compatible, scalable DSP software model • Designed for multi-core implementations
Differentiation	• Only include the acceleration options needed for implementation • Customize further with TIE if needed
Virtually unlimited bandwidth	• Add GPIO, FIFO, or SRAM interfaces up to 1024 bits each • Create predictable connections to/from the rest of the SoC
Functional safety	• ISO26262 SEooC for ASIL-B and ASIL-D products

For more information, please visit https://ip.cadence.com

Tensilica Product Portfolio

Configurable Multi-Core DSP and AI Platform

Audio/Voice

- Audio post-processing
- Voice pre-processing
- Speech AI/Enhancement
- Voice trigger/recognition
- Audio encode and decode

Radar/Lidar/Comms

- Automotive Radar/Lidar
- Always-alert Sensor
- V2X, 5G, LTE, Wireless
- WiFi, Smart Grid
- Infrastructure and Terminals

Vision

- Image/Vision Processing
- Low end AI at the Edge
- AR/VR
- ADAS

FloatingPoint

- AI/ML
- Motor Control
- Sensor Fusion
- Object Tracking
- AR/VR
- High Performance Computing

AI Processor

- Stand Alone AI processor
- AI at the Edge
- CNN, RNN, regression

Xtensa® CPU

- Pre-configured
- Low Power
- Deeply Embedded
- DSP + Control

Xtensa®: The Configurable, Extensible Processor Platform

Custom ISA, Accelerators, Data Types, and Interfaces (TIE: Tensilica Instruction Extension)

ConnX DSPs

ConnX Advantages

- Industry leading radar, lidar, and V2X DSPs
- Certified compliant for ISO26262:2018 **ASIL-B or ASIL-D**
- Optimized complex data processing (fixed and float)
- Linear algebra acceleration
- Performance-enhanced FFT, FIR computation
- Optimized for CFAR, beamforming, convolution processing
- Accelerations for communications (Viterbi, Turbo, LDPC)
- Optimized NatureDSP library, application SW examples

Main Functions

- Radar sensing
- Lidar sensing
- V2X
- Gesture sensing
- Sensor fusion
- Baseband PHY Layer
- Beamforming

Eigen Library for ConnX DSPs

- High-level library for linear algebra, matrix, vector, and numerical solvers
- High-performance, reliable, and widely used to speed up math routines
  - Explicit vectorization is performed for SIMD instruction sets
- Well-known API and industry standard
  - Used by CPU, GPU, and numerous platforms
  - The API is clean and expressive while feeling natural to C++ programmers
  - Eigen library can be downloaded from http://eigen.tuxfamily.org/

$$AX = B$$
$$A^{-1}(AX) = A^{-1}B$$
$$(A^{-1}A)X = A^{-1}B$$
$$(I_n)X = A^{-1}B$$
$$X = A^{-1}B$$

- Cadence provides the following components that can be overwritten over the downloaded open-source Eigen library

Component	Description
PacketMath.h and related header files	DSP architecture-specific header files for helping Eigen library to do better vectorization. It includes baseops and math functions. Details of currently supported data types showed in the table below.
Eigen library CMake related files	The functionality and performance test files updated for Xtensa compiler and ISS
Eigen library support for full and half packet enabled	Eigen library provides a way to add support of half packet (half-SIMD) processing to improve performance of operations on non-SIMD multiple matrices. This half-packet support is added in *PacketMath*.

ConnX DSP PacketMath Support

	Single	Double
real	✓	✗
complex	✓	✗

Double-precision data types will be supported soon.

NatureDSP Libraries for ConnX DSPs

Rich and optimized DSP libraries for basic and advanced DSP functions

- Scalar/Vector fixed/floating-point, real/complex support
- Library source code available

Note: To cover the complete library routine list, all optional configurations have to be selected

Category	Functions
Math	Scalar and vector for: arc cosine/sine/tangent, cosine/sine/tangent/cotangent, hyperbolic sine/cosine/tangent, sigmoid, logarithm, anti-logarithm, exponential, square root, reciprocal, reciprocal square root, softmax
Complex	Magnitude, phase, conjugate, exponent, combined cosine and sine, normalization, division, polar to cartesian and cartesian to polar conversion
FIR Filters	Convolution, auto/cross correlation, interpolation, decimation, polynomial fitting/interpolation
IIR Filters	Biquad, lattice block IIR
FFT	Complex, real, FFT/IFFT, DFT, blockwise, and streaming
Vector Operations	Product, sum, magnitude, reciprocal, division, transcendental, peak search, mean
Matrix Operations	Multiply, Transpose, Hermitian Product, Streaming-Block order conversion, Interleaving/De-interleaving
Matrix Decomposition and Inversion	QR decomposition and Back substitution, Cholesky MMSE Solver, Cholesky Decomposition/Forward and Backward substitution, Direct Matrix Inversion, Matrix Inversion by Gauss-Jordan, LU Decomposition, Determinant, Linear equation solution, Singular Value Decomposition, Eigenvalues and Eigen vectors
Communications	CRC, de-spreading, modulation, slicer, convolutional encoding, bit manipulation, PRBS Generation, space time coding

Radar Libraries for ConnX DSPs

Optimized DSP functions for Radar Applications

- Scalar/Vector fixed/floating-point, real/complex support
- Library source code available

Note: To cover the complete library routine list, all optional configurations have to be selected

Category	Functions
Range FFT	Range FFT(RFFT) is used to calculate range of the targets by computing forward Fast Fourier Transform(FFT) on real/complex input data with/without window. It calculates blockwise FFT on L blocks of N input samples, where supported values of N (FFT size) are the power of 2 (2 ^ n) where n = 4, 5, - - - 12 and L>0.
Doppler FFT	This function computes the forward FFT on complex streaming order input data with/without windowing. It provides three types of outputs which are complex FFT spectrum, real normalized magnitudes of complex FFT spectrum, and Log (base 2) of normalized magnitudes of FFT spectrum.
CFAR	False Alarm Rate (CFAR) module is an adaptive thresholding technique that detects targets present in a Range Doppler Image (RDI). Each pixel of the RDI is called a cell, and whether the target is present in the cell or not is determined by CFAR algorithm. A specific cell for which this algorithm is applied is referred to as the Cell Under Test (CUT).
Tracking: Extended Kalman Filter	Kalman filter addresses the general problem of estimating the state of a discrete-time controlled process governed by a linear stochastic difference equation. If the process to be estimated and (or) the measurement relationship to the process is non-linear, we can linearize the estimation around the current estimate using the partial derivatives of the process and measurement functions to compute estimates. The Extended Kalman Filter (EKF) is the nonlinear version of the Kalman filter, which linearizes an estimate of the current mean and covariance

Cadence Tensilica Code Generation Toolbox

- Matlab toolbox-based package is provided for ConnX B10/B20 DSPs
  - User guide describes the details of the installation procedure
- Package provides both Simulink and Matlab code generation support
  - Generate code targeting the Cadence Tensilica ConnX DSPs
    - Code Replacement for certain kernels from NatureDSP library is provided
    - NatureDSP library is required to be built separately and provided during the setup
  - Compile optimized code using the Cadence Xtensa C/C++ Compilers
  - Run the compiled code using the Cadence Xtensa Instruction Set Simulator
  - Check the results and behavior of code running in Processor-In-the-Loop (PIL) manner against the Matlab/Simulink's simulated test vectors.
- Windows and Linux environments are supported
  - Current package works with Matlab R2021a
- Toolboxes required
  - Embedded Coder
  - Matlab Coder
  - Simulink and Simulink Coder (Required for Simulink code generation)

Industry Leading Radar / Lidar DSPs for Automotive Applications

Autonomous Vehicle Design Win

- Global Tier 1 Semiconductor Vendor
- Autonomous Vehicle
- Multiple ConnX B20 DSPs
- 4D Imaging Radar

Radar Design Wins Globally

- NA, EMEA, APAC
- Global Tier 1 Semiconductor Vendors
- RTL Delivered
- Silicon in Mass Production

Tensilica ConnX - Scalable Radar / Lidar Processing Solution

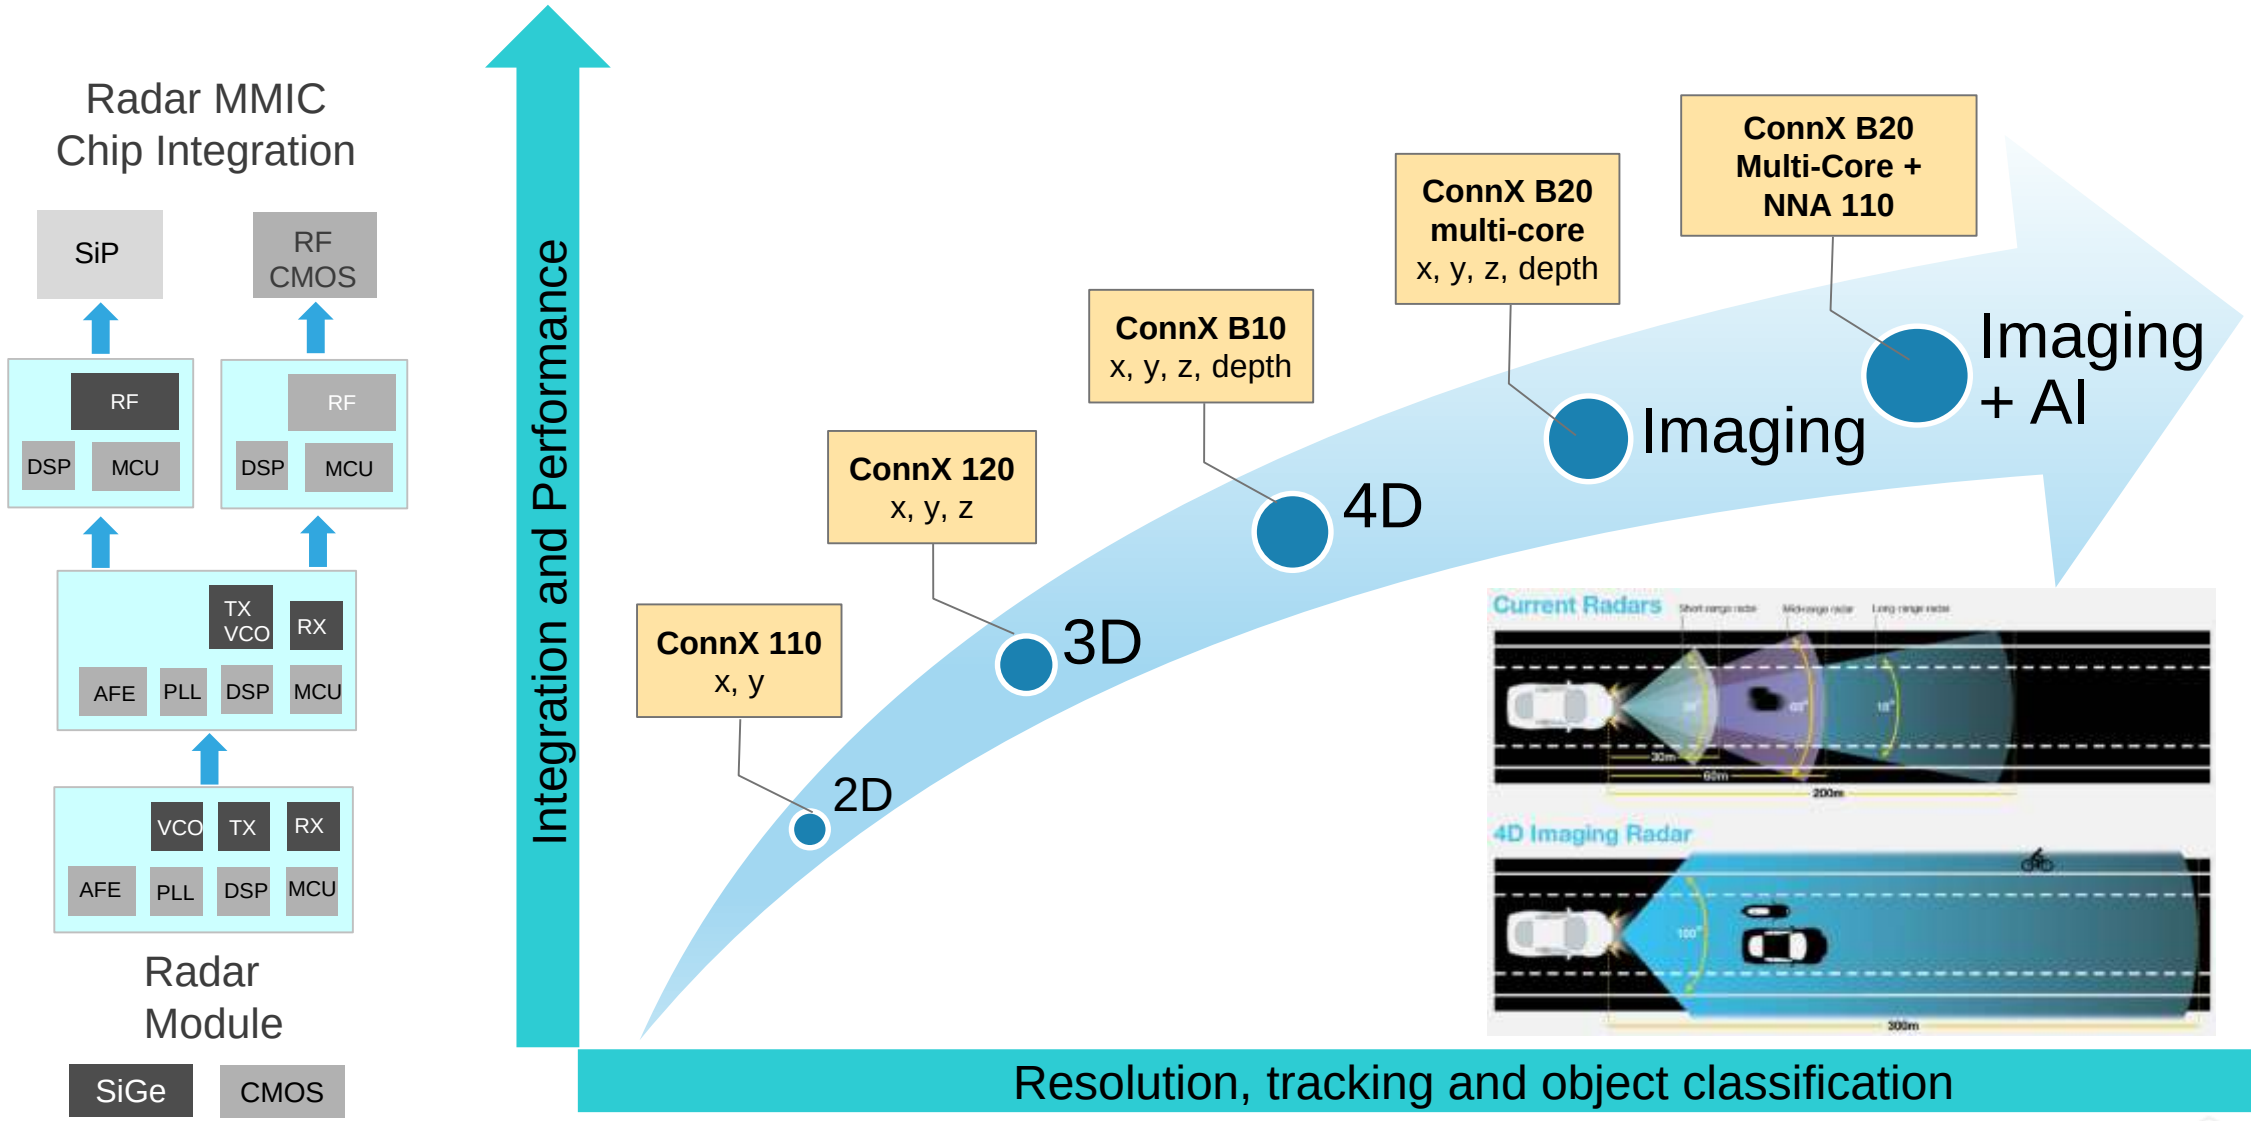

CEVA®

PentaG-RAN™
by CEVA

PentaG-RAN
5G RAN Baseband Platform IP for Cellular Infrastructure Solutions

Application-Specific Processors
November 1, 2022

Eric Cowden
FAE

www.ceva-dsp.com

Technology Portfolio – Enabling Smart Edge Devices

Wireless Connectivity IPs	Smart Sensing IPs

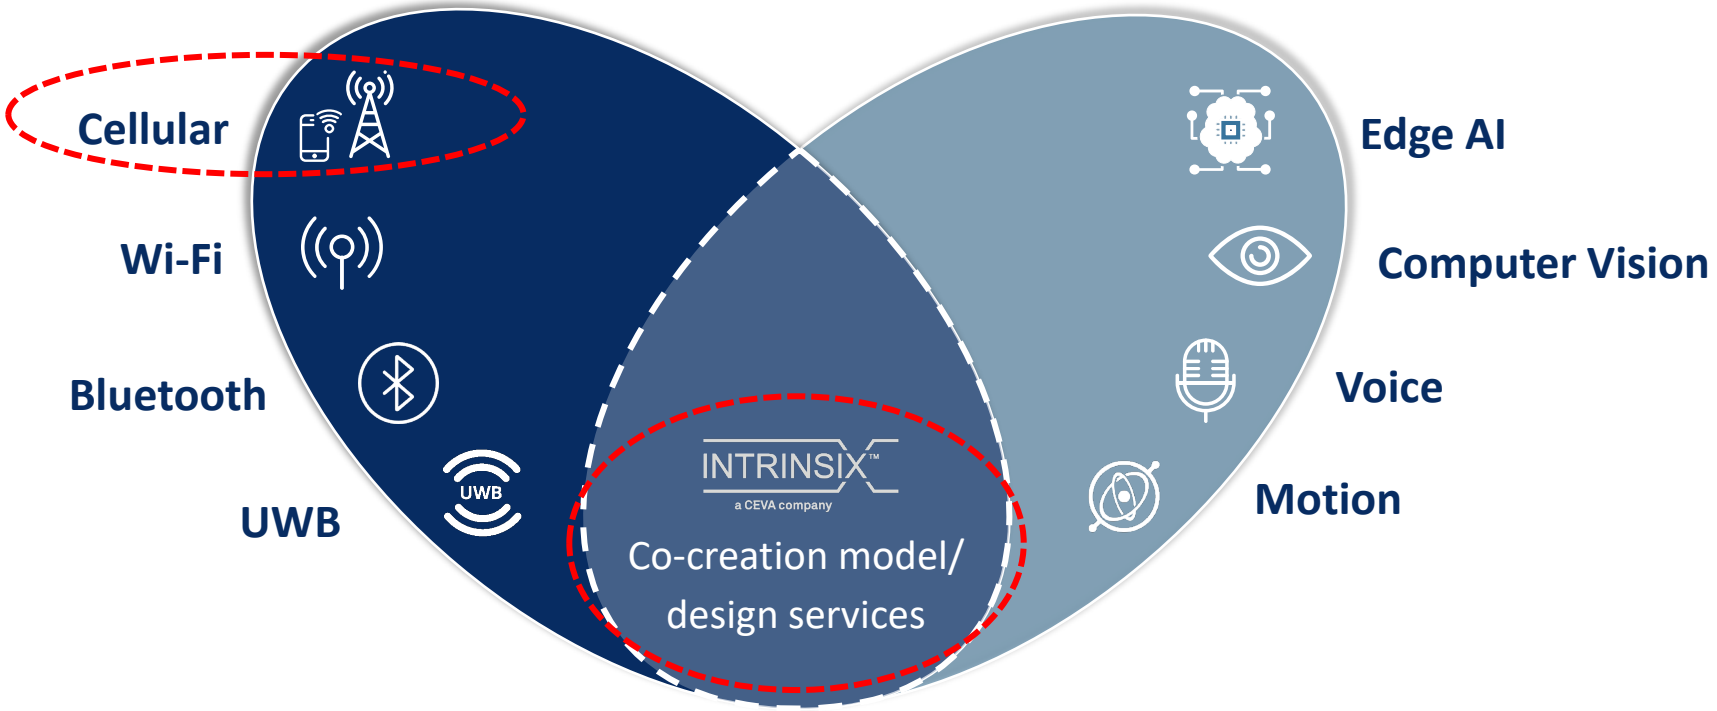

Building on CEVA's Legacy as Leading Cellular IP Vendor

From Infrastructure to IoT

>15m 5G base station SoCs

5G RAN

- **CEVA's technology extends to all 5G RAN architectures**
  - PentaG-RAN™ by CEVA - New
- Base-station, disaggregated DU/RU, AAU, small cells, vRAN, OpenRAN backhaul, fronthaul

>10bn handsets

Smartphones

- **Multiple technologies to address smartphones in 5G era**
- Baseband modem – **PentaG2**
  Bluetooth / Wi-Fi / UWB

>100m cellular IoT modems

Cellular IoT

- **CEVA offers DSPs & platforms for wide range of Cellular IoT technologies**
- NB-IoT, RedCap, Cellular V2X
  Industrial IoT, FWA, 5G Satellite

NOKIA ZTE UNISOC intel SAMSUNG NORDIC SEMICONDUCTOR ASR MICROELECTRONICS SEQUANS Autotalks

5G RAN Technology Trends

- 5G networks are evolving in several vectors, all pointing towards network openness and sophistication:

Massive MIMO and Beamforming (BF)

- Radio Unit (RRU) is becoming a Massive MIMO 32/64/128 antenna Active Antenna Unit (AAU) with massive compute requirements
- Significant complexity in beamforming computation and management – need for baseband processing in RRU in antenna site

Open RAN (O-RAN)

- O-RAN specifying common and open interfaces between RAN components
- Emergence of new OEMs **– need for new merchant silicon with efficient baseband processing** to replace power and cost inefficient FPGAs and COTS platforms

Virtualized RAN (V-RAN)

- VRAN - Running virtual SW defined networks on COTS platforms
- Aggregation of multiple links on single platform - Massive compute requirements for Baseband Unit

Proliferation of Small Cells

- For indoor, dense deployments, and mmWave
- Large volume and easier penetration – new players

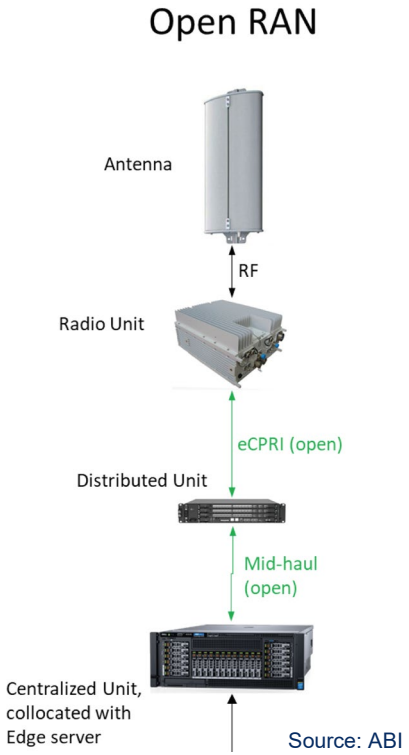

5G RAN Massive MIMO Technology Challenge

- Immense compute needs for massive 64x64 antenna arrays. Existing COTS solutions are not efficient and cannot scale to massive MIMO dimensions
- Diverse workloads in 5G physical layer require complex and heterogeneous L1 sub-systems, and optimal HW/SW partitioning
- New Entrants to the market are intimidated by the challenges posed to develop ASICs for this huge market opportunity
  - There is a design gap within the semiconductor industry for 5G baseband processing, where there is a scarcity of design and architecture expertise, particularly within the PHY and radio domains, which are highly complex in nature

The Need for Cost Effective ASIC Solutions

On DU Side

- Get much more value: accelerate complete L1 processing chain vs FEC only acceleration

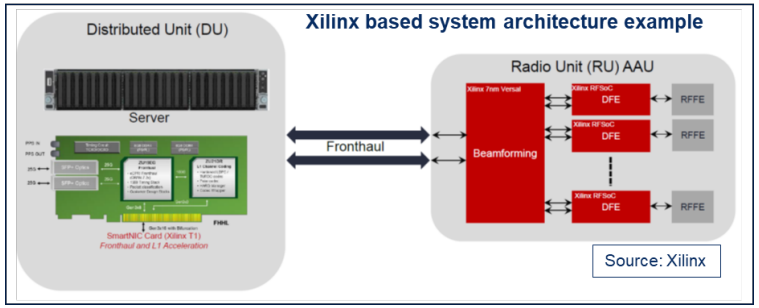

10X lower power and area with PentaG-RAN

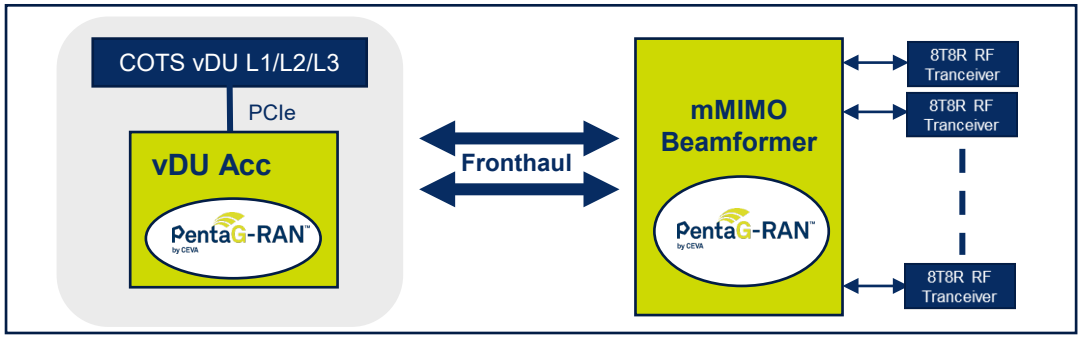

On RU Side

- Replace costly FPGAs with single efficient silicon
- Various integration options with transceiver

PentaG-RAN™ by CEVA

Industry's First 5G Baseband Platform IP for Open RAN ASICs

- **PentaG-RAN** is targeting cellular infrastructure ASICs in **base station** and **radio** configurations, from small cell to Massive MIMO
- **PentaG-RAN** is a heterogeneous baseband compute platform that provides a complete licensable L1 PHY solution with optimal hardware/software partitioning, integrating best in class vector DSPs and complete acceleration of signal datapath using efficient hardware accelerators
- **PentaG-RAN** platform delivers up to **10X savings** in power and area compared to available FPGA and COTS CPU based alternatives
- **PentaG-RAN** is also offered in SoC co-creation design model using CEVA-Intrinsix

PentaG-RAN
Configuration for Small Cell and DU

PentaG-RAN 5G DU/Small-Cell Platform

PentaG-RAN™ by CEVA

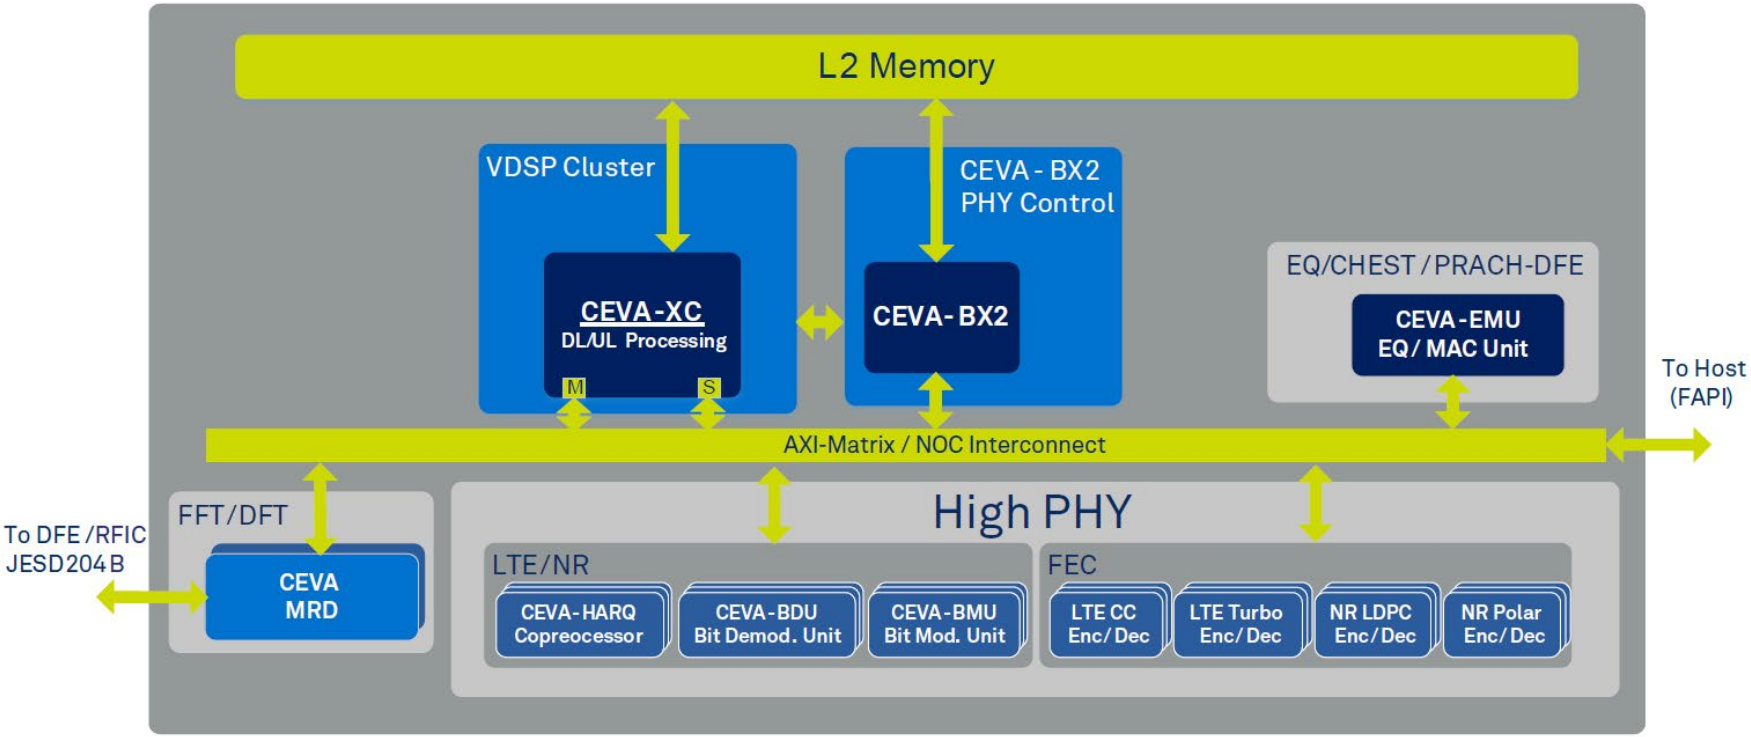

Optimal HW/SW Partitioning

PentaG-RAN™ by CEVA

- Complete acceleration of signal path processing. Wide range of acceleration blocks
- Flexible SDR platform for running side-algorithms (e.g. channel estimation) allowing customer differentiation and secret sauce

5G UL Processing Chain Data Path Acceleration

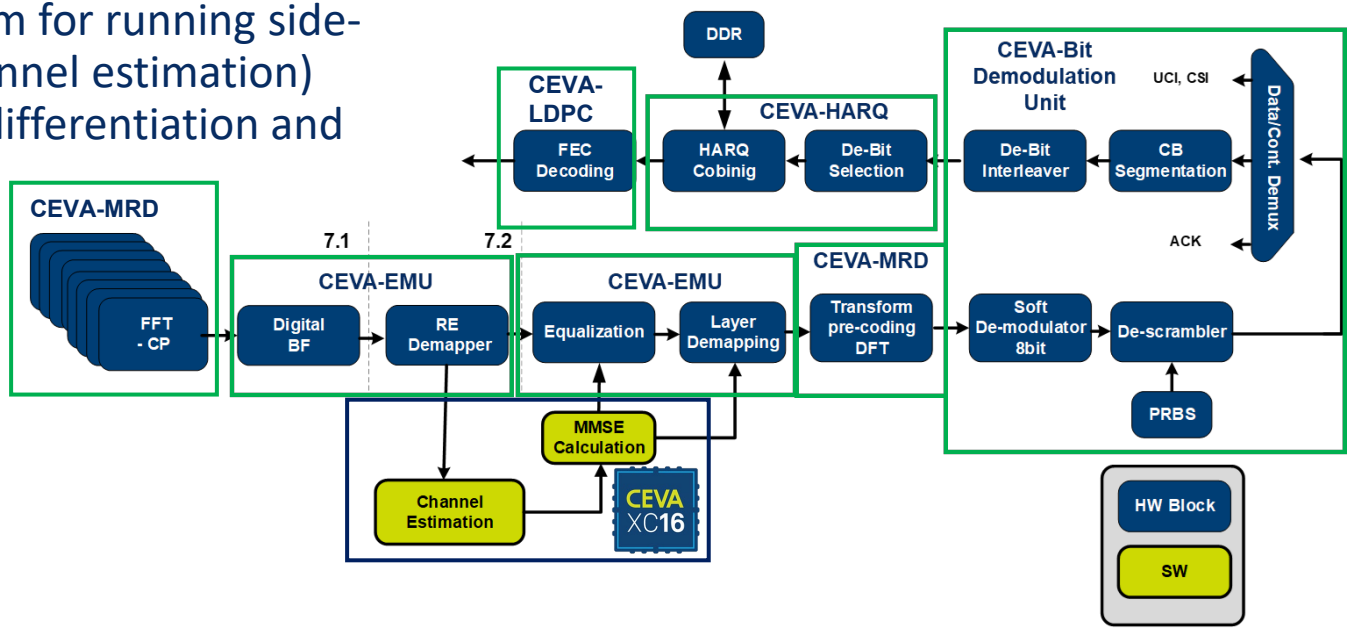

Massive vRAN DU Acceleration

PentaG-RAN™ by CEVA

- PentaG-RAN can be used as vRAN inline acceleration compute resource to COTS platforms (x86 or ARM based)
- Based on market leading CEVA-XC
- Baseband platform also includes:
  - CEVA-BX2 scalar controller/DSPs for PHY control and sequencing
  - **Comprehensive** set of HW accelerators and co-processors for accelerating **complete** data path

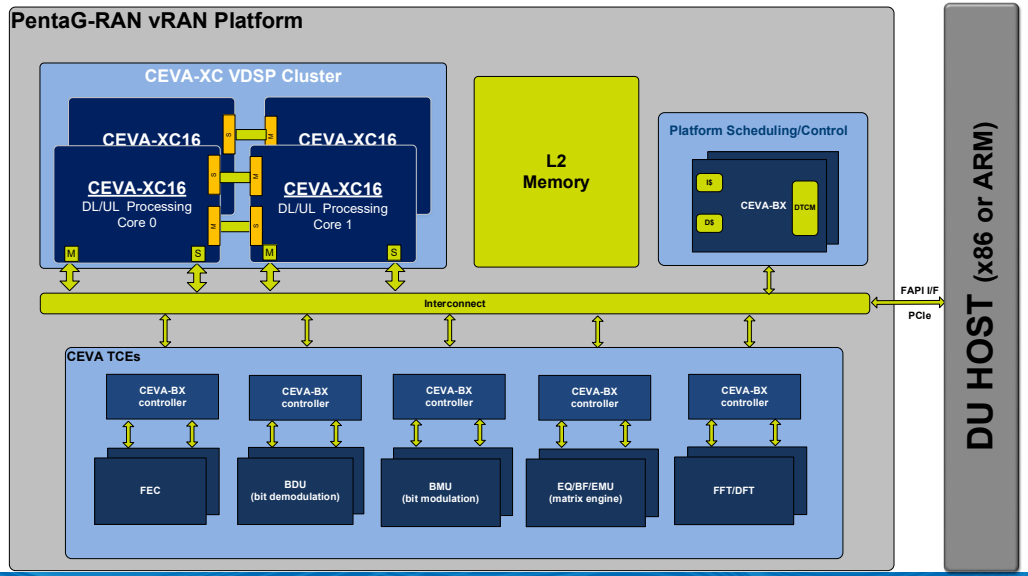

PentaG-RAN
Configuration for RRU and Massive MIMO

Typical Open RAN Radio Unit

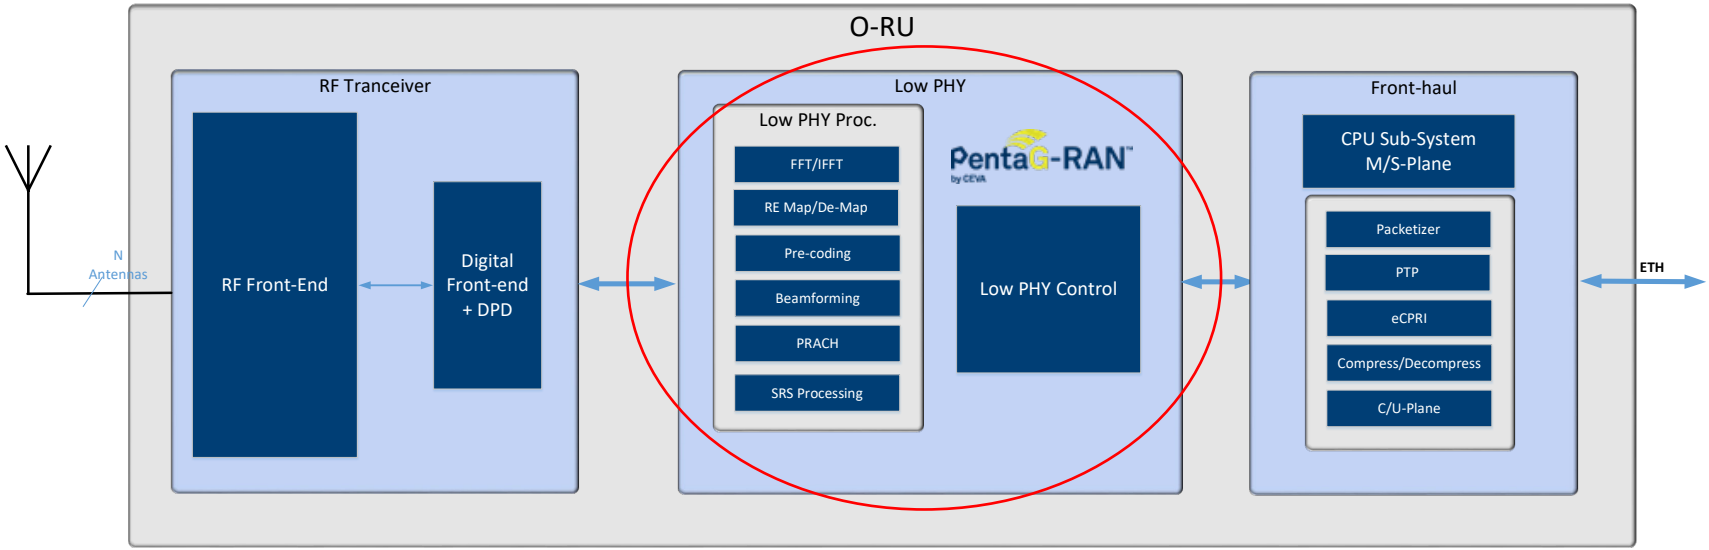

Massive MIMO Beamformer SOC

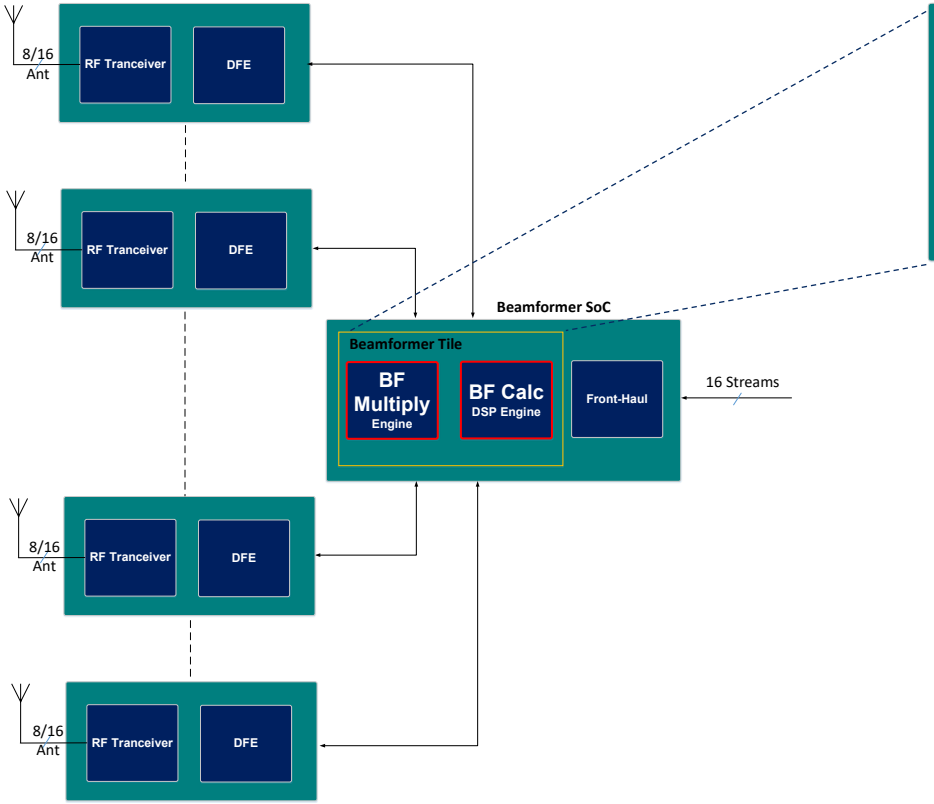

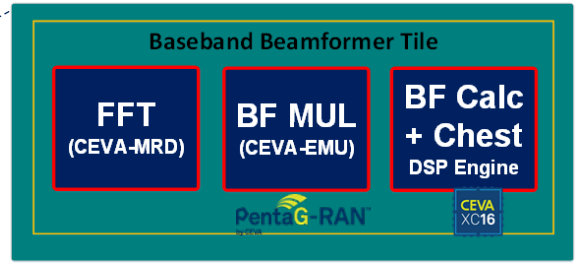

PentaG-RAN Digital Beamforming:
- BF Weight Compute based on CEVA-XC DSP
- CEVA-EMU BF Matrix Multiply Engine
- CEVA-MRD FFT Engine

Beamformer sub-system/Tile can be embedded within central Beamformer SoC or transceiver chip/SiP
- Option for scalable chiplet architecture

PentaG-RAN Massive MIMO Radio Platform

PentaG-RAN™ by CEVA

- Heterogeneous scalable IP platform comprised from DSP cores (best in class XC16 and BX2) and HW accelerators (co-processors)
- Optimal HW/SW partitioning allowing solution scaling and meeting competitive power consumption targets
- DSPs handle advanced algorithms like channel estimation, BF calculation, etc.
- Data path fully mapped to HW acceleration

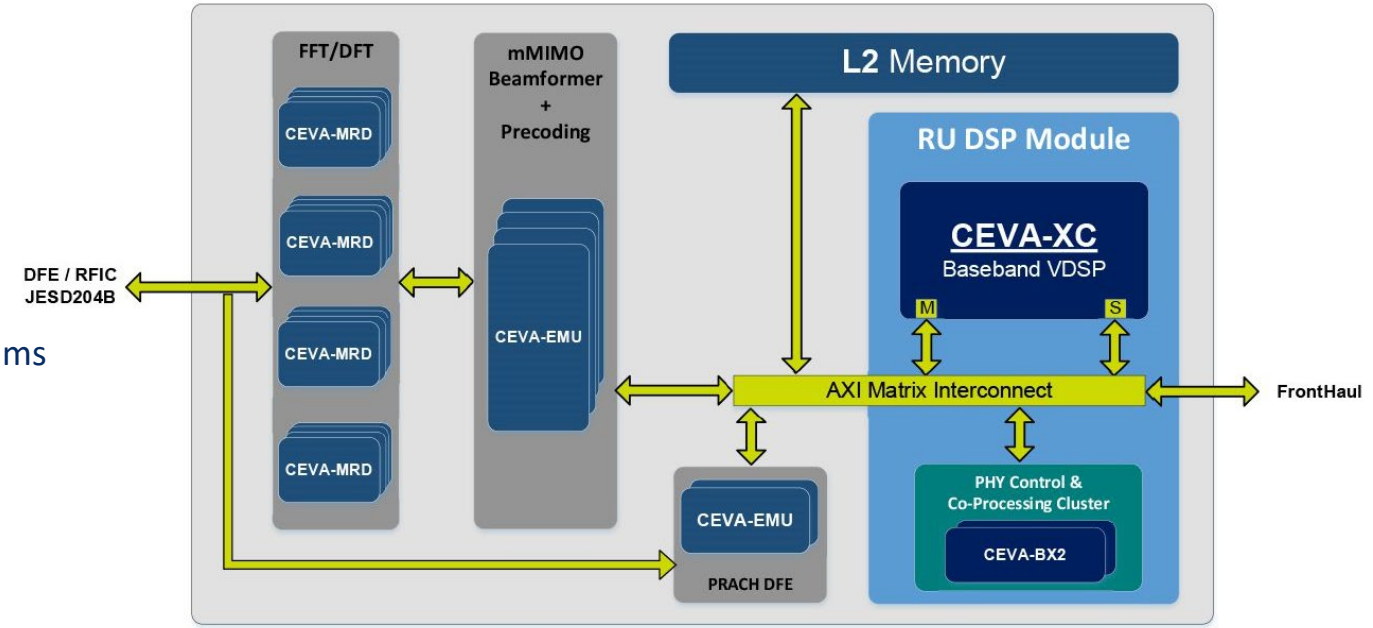

Virtual Platform Simulator

- An integrated System-C simulation environment that allows system engineers, architects and SW developers to model, profile and debug in pre-silicon
- Supporting CEVA's vector & scalar DSP cores and HW accelerators
- Fast prototyping and architecture definition
- Strong and effective IP evaluation and PoC tool
- Seamless MATLAB interface for algorithmic development and tests
- Available for evaluation and demos

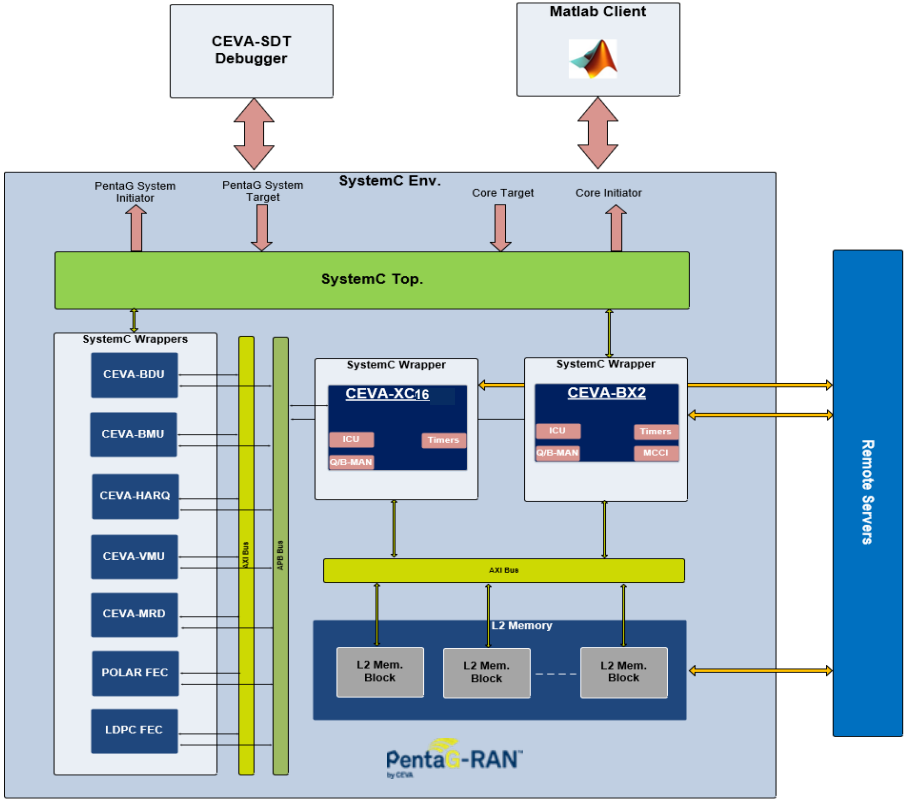

Trusted Partnership: SoC Co-Creation

CEVA's 5G Co-Creation Offering

- CEVA's PentaG-RAN is a unique 5G platform IP offering allowing customers to quickly introduce their own 5G SoCs
- CEVA can customize this comprehensive solution to fit Customer's needs
  - Map customer use case to PentaG-RAN platform
  - Complete architecture spec, solution dimensioning
  - Interconnect definition (to transceiver, JESD, Fronthaul, eCPRI, etc.)
  - Adaptation to specific process node and operating conditions
  - SW architecture, and foundation SW framework

- Intrinsix team (SoC design services of CEVA) can supply complete modem design services, from architecture up to GDS

➔ **These steps ensure Customer can identify its differentiation opportunity as well as offload development resources**

DU/Small Cell SoC Design Co-creation Example

- Requires complex interconnect options and potentially L2/L3 CPU cluster
- CEVA, with its Intrinsix team, offers co-creation services from L1 architecture, to full L1 design service, to complete SoC design, up to GDS

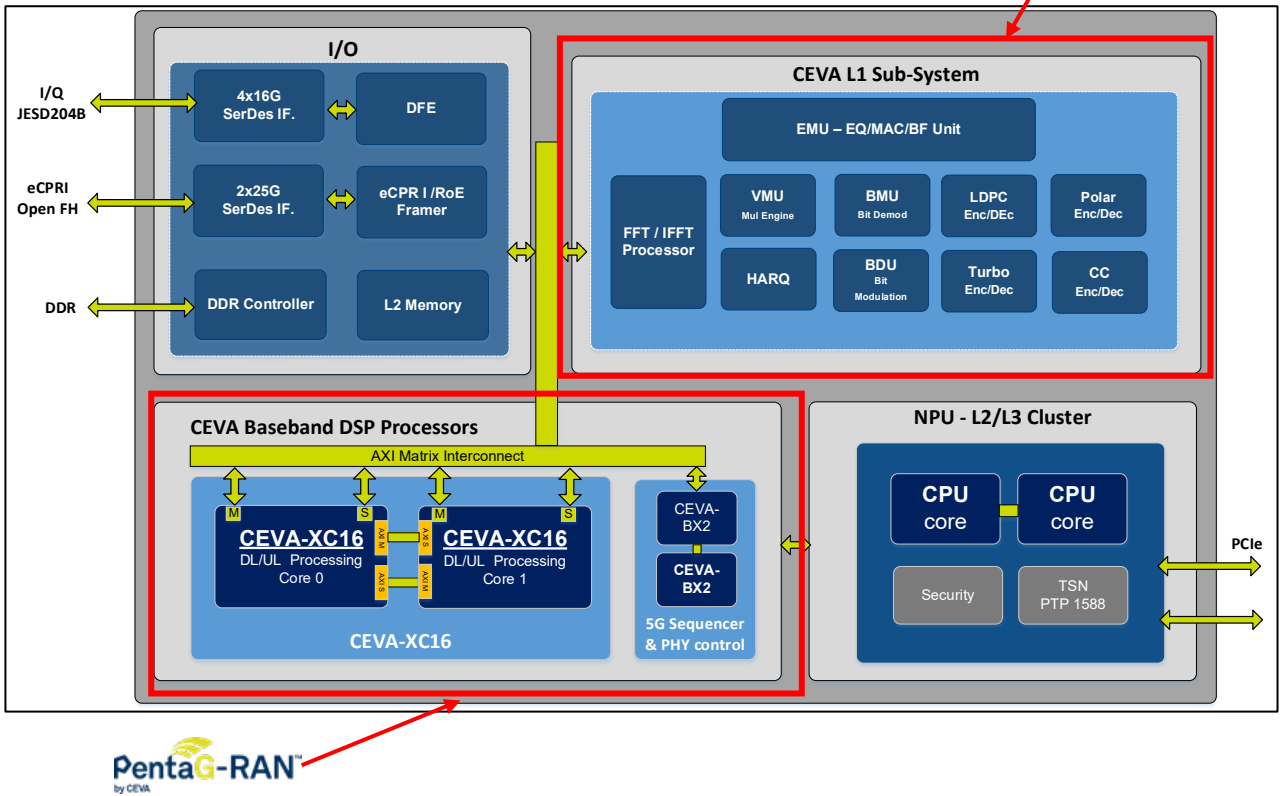

Thank You

www.ceva-dsp.com

MARVELL

DPUs: Where Will They Go Next?

Manoj Roge
Sr Director, Processor Business Unit, Marvell

Linley Fall Processor Conference, November 2022

Forward-looking statements

Except for statements of historical fact, this presentation contains forward-looking statements (within the meaning of the federal securities laws) including, but not limited to, statements related to market trends and to the company's business and operations, business opportunities, growth strategy and expectations, and financial targets and plans, that involve risks and uncertainties. Words such as "anticipates," "expects," "intends," "plans," "projects," "believes," "seeks," "estimates," "can," "may," "will," "would" and similar expressions identify such forward-looking statements. These statements are not guarantees of results and should not be considered as an indication of future activity or future performance. Actual events or results may differ materially from those described in this presentation due to a number of risks and uncertainties.

For factors that could cause Marvell's results to vary from expectations, please see the risk factors identified in Marvell's Quarterly Report on Form 10-Q for the fiscal quarter ended July 30, 2022, as filed with the SEC on August 26, 2022, and Marvell's Annual Report on Form 10-K for the fiscal year ended January 29, 2022, as filed with the SEC on March 10, 2022, and other factors detailed from time to time in Marvell's filings with the SEC. The forward-looking statements in this presentation speak only as of the date of this presentation and Marvell undertakes no obligation to revise or update publicly any forward-looking statements.

Agenda

1. System trends
2. OCTEON®: The original DPU platform
3. DPU use cases
4. OCTEON® solutions & benchmarks

DPU need driven by data centric applications

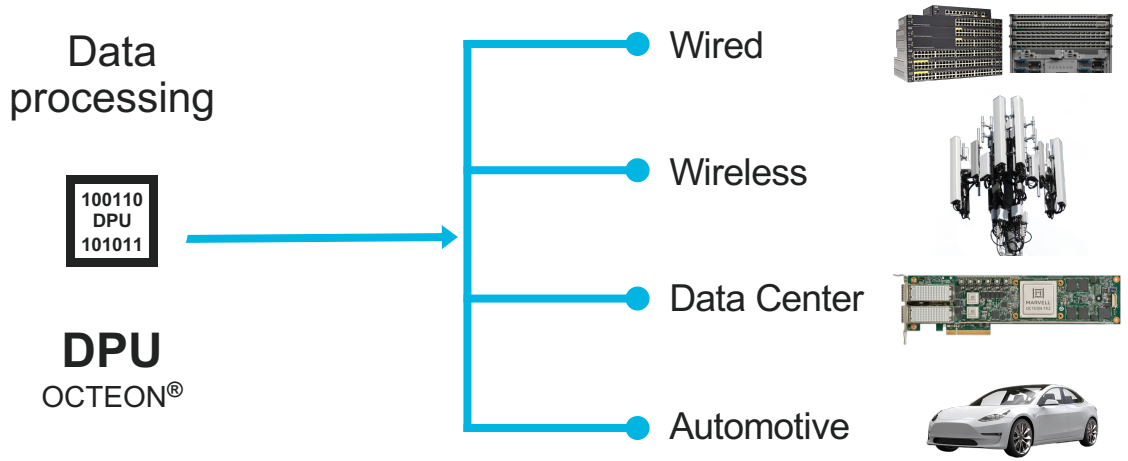

Multiple end markets for
OCTEON® DPU

Future Data Centers:
Composable, software-defined, hardware-accelerated

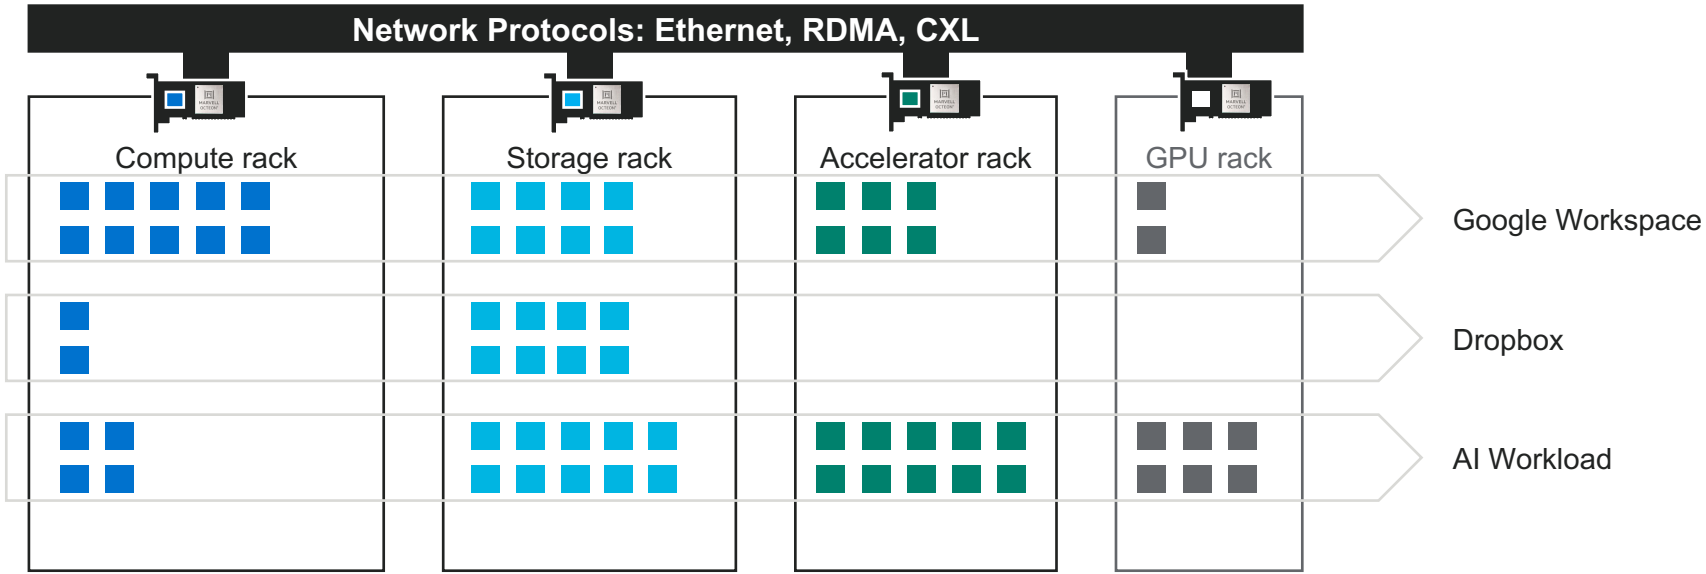

DPUs manage composability, accelerate workloads

DPU in every server!

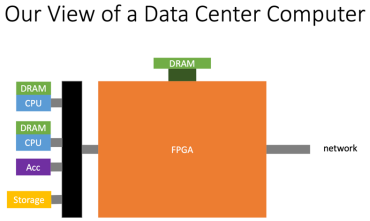

Microsoft

H2RC'16: "CPU is complexity offload engine for FPGA!"[1]

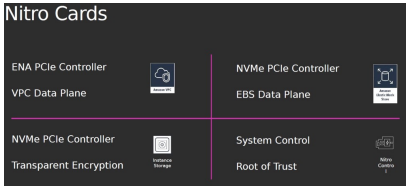

Amazon

2019 re:Invent – "All new instance launches use the Nitro System"[2]

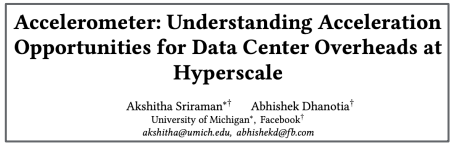

Accelerometer: Understanding Acceleration Opportunities for Data Center Overheads at Hyperscale

Akshitha Sriraman*† Abhishek Dhanotia†
University of Michigan*, Facebook†
akshitha@umich.edu, abhishekd@fb.com

Meta

"Microservices spend as few as 18% of CPU cycles executing core application logic"[3]

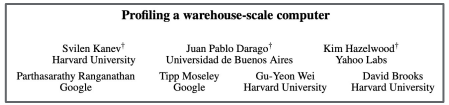

Profiling a warehouse-scale computer

Svilen Kanev† Juan Pablo Darago† Kim Hazelwood†
Harvard University Universidad de Buenos Aires Yahoo Labs
Parthasarathy Ranganathan Tipp Moseley Gu-Yeon Wei David Brooks
Google Google Harvard University Harvard University

Google

"Data center Tax" can comprise nearly 30% of cycles[4]

1: H2RC 2016 keynote.
2: AWS reinvent 2019
3: Accelerometer paper
4: Profiling a warehouse-scale computer paper

Transition to virtualization

Traditional wireline appliances

Data center

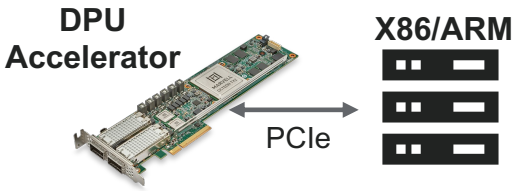

Traditional RAN/carrier

Data center

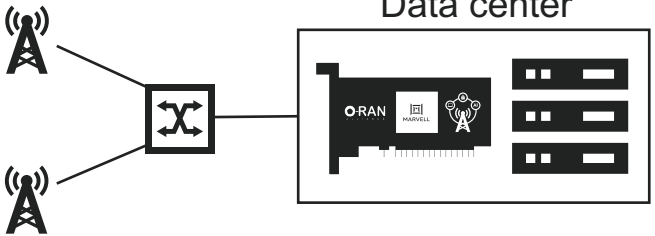

vRAN

Applications require versatile mix of accelerators

Data center

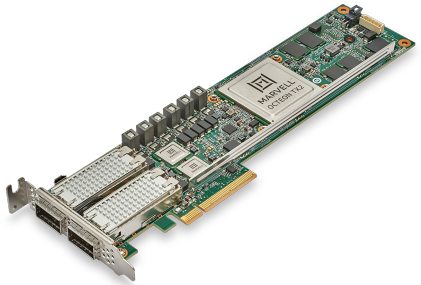

Networking Storage Security AI/ML

5G Infrastructure

Baseband Security AI/ML Networking

OCTEON®: The original DPU platform

CAVIUM

OCTEON® multicore

2005

Industry's 1st DPU

LiquidIO®

2010

OCTEON® Fusion

2012

OCTEON® TX

2015

First Marvell Arm DPU

OCTEON® TX2

2019

5th Generation Marvell DPU

OCTEON® 10

2021

6th Generation Marvell DPU

OCTEON® Next Gen

Enterprise

Cloud

Wireless

Radio

Edge

Auto

OCTEON® 10 architectural overview

Scalable Compute

- ARMv9.0 64-bit Neoverse N2 cores

Memory subsystem and connectivity

- 1MB/core L2, 2MB shared last level cache
- DDR5 w/ sideband-ECC and memory encryption
- XCalibur mesh interconnect

Hardware acceleration

- Highly-virtualized, software-friendly NIC
- Packet processing, QoS, hierarchical queues with shaper and WDRR scheduler
- Inline and Co-processor security (SSL/IPSec)
- Compression, Decompression
- Inline ML inference engine
- Secure boot + embedded hardware security module

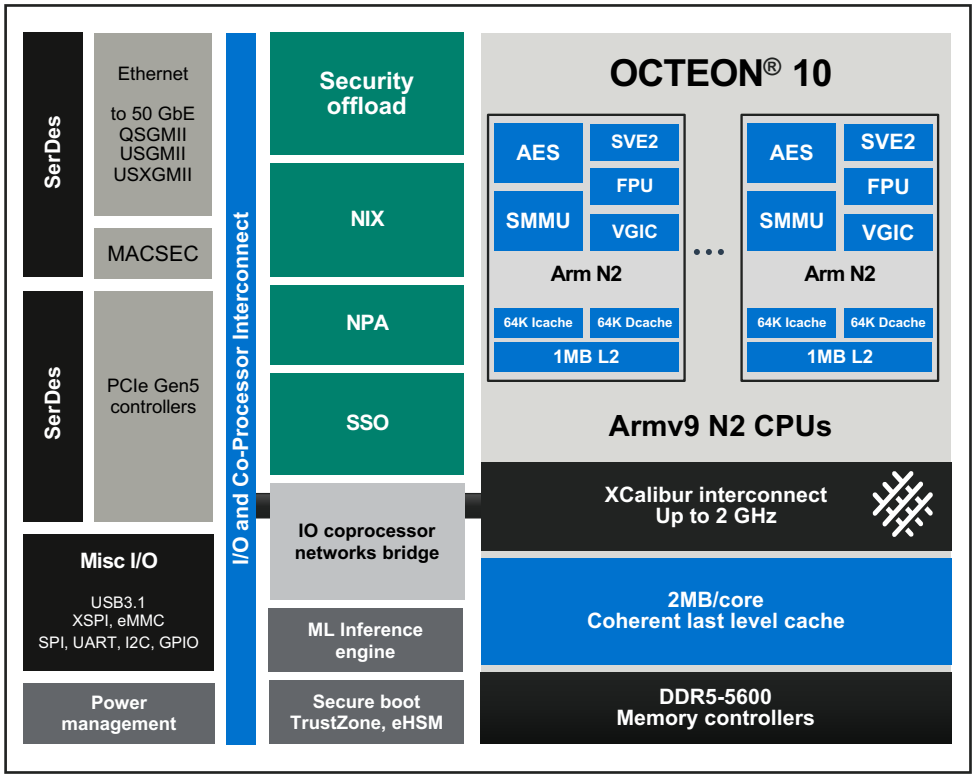

Platform strategy

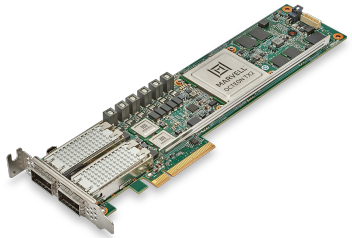

+

Virtualization

APIs and frameworks

Apps and accelerators

- DPU PCIe cards
- Robust and open source software support
- Partner ecosystem

Marvell PCIe accelerator cards

	CN98	CN106	CN103
Part Number	WA-CN98-A1-PCIE-4P100-R1	WA-CN106-A1-PCIE-2P100-R1	WA-CN103-A0-PCIE-4P50-R1
Port config	4x 100G PAM4	2x 100G PAM4	4x 50G PAM4
PCIe	Gen4	Gen5	Gen5
Core	36x ARM V8 TX2	24x ARM V9 N2	8x ARM V9 N2
Availability	Now		1Q CY23

Announcing
General availability!

CN106 based PCIe card

Part number	WA-CN106-A1-PCIE-2P100-R1
Features	**Capability**
I/O	2 x 100G PAM4 PCIe Gen 5
Memory	6 x 40bit DDR5@ 5200MTs w/ECC, 8-40GB total
ARM cores	24 ARM N2, 2.5GHz, 100 SPECINT2017
Performance	120 MPPS, 120Gbps
IPSEC, RSA 2K, 1KB OpenSSL, TLS1.3 support	120Gbps IPSEC, 24Kops RSA 2K, 120Gbps 1KB OpenSSL
Hard ML block	Yes, 16TOPS

Availability	Date
SDK11 support	Now
Order in qty	Nov 2022

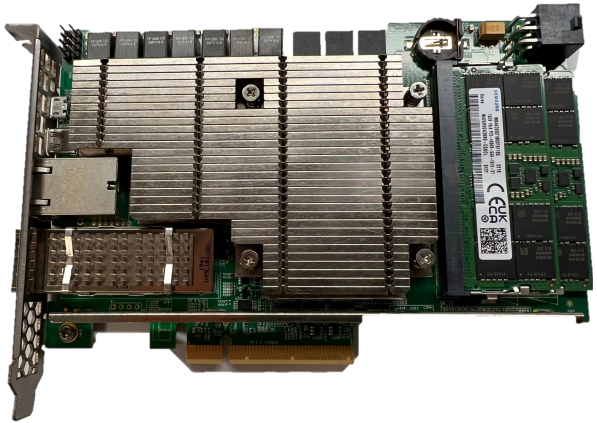

DPU solutions

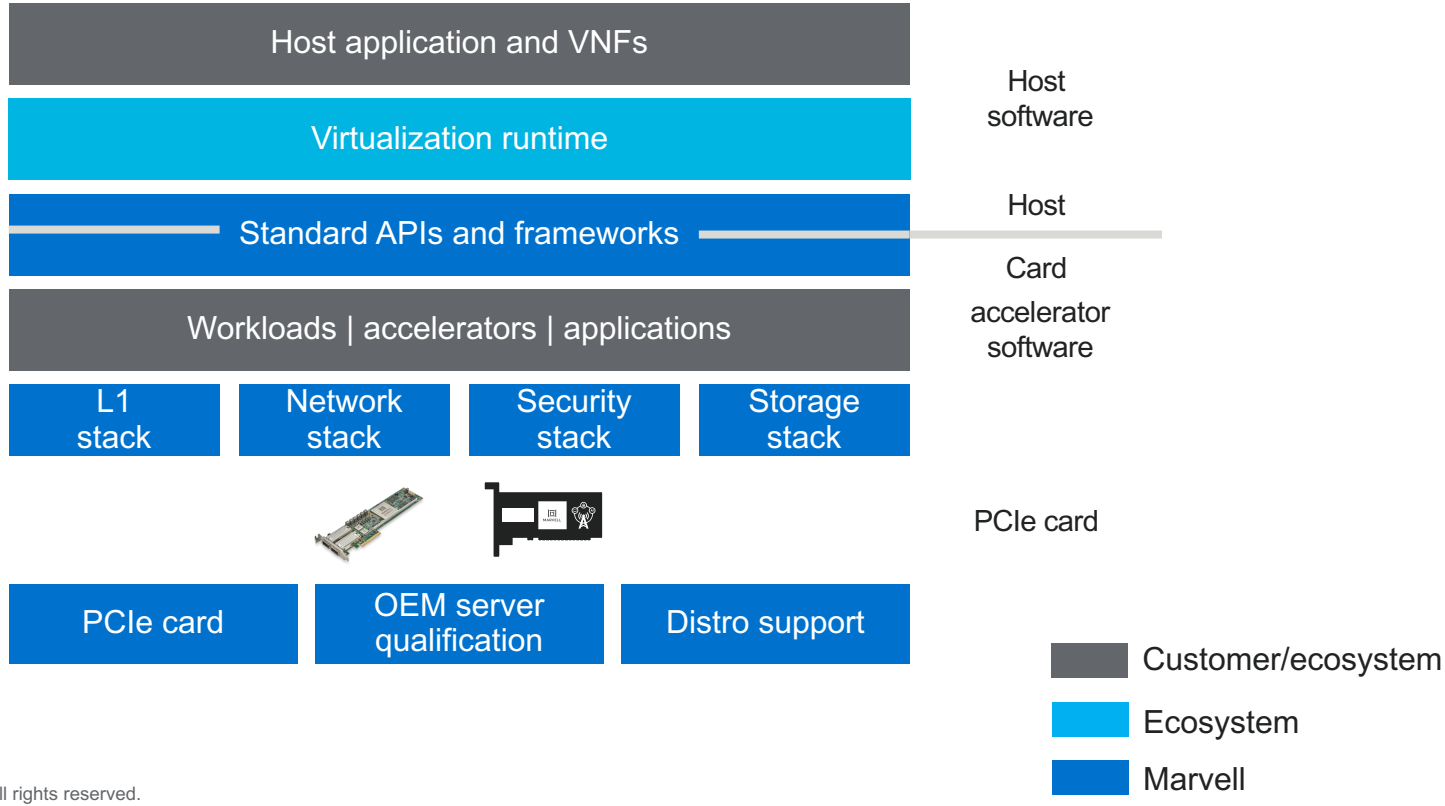

Industry-leading 5G infrastructure portfolio

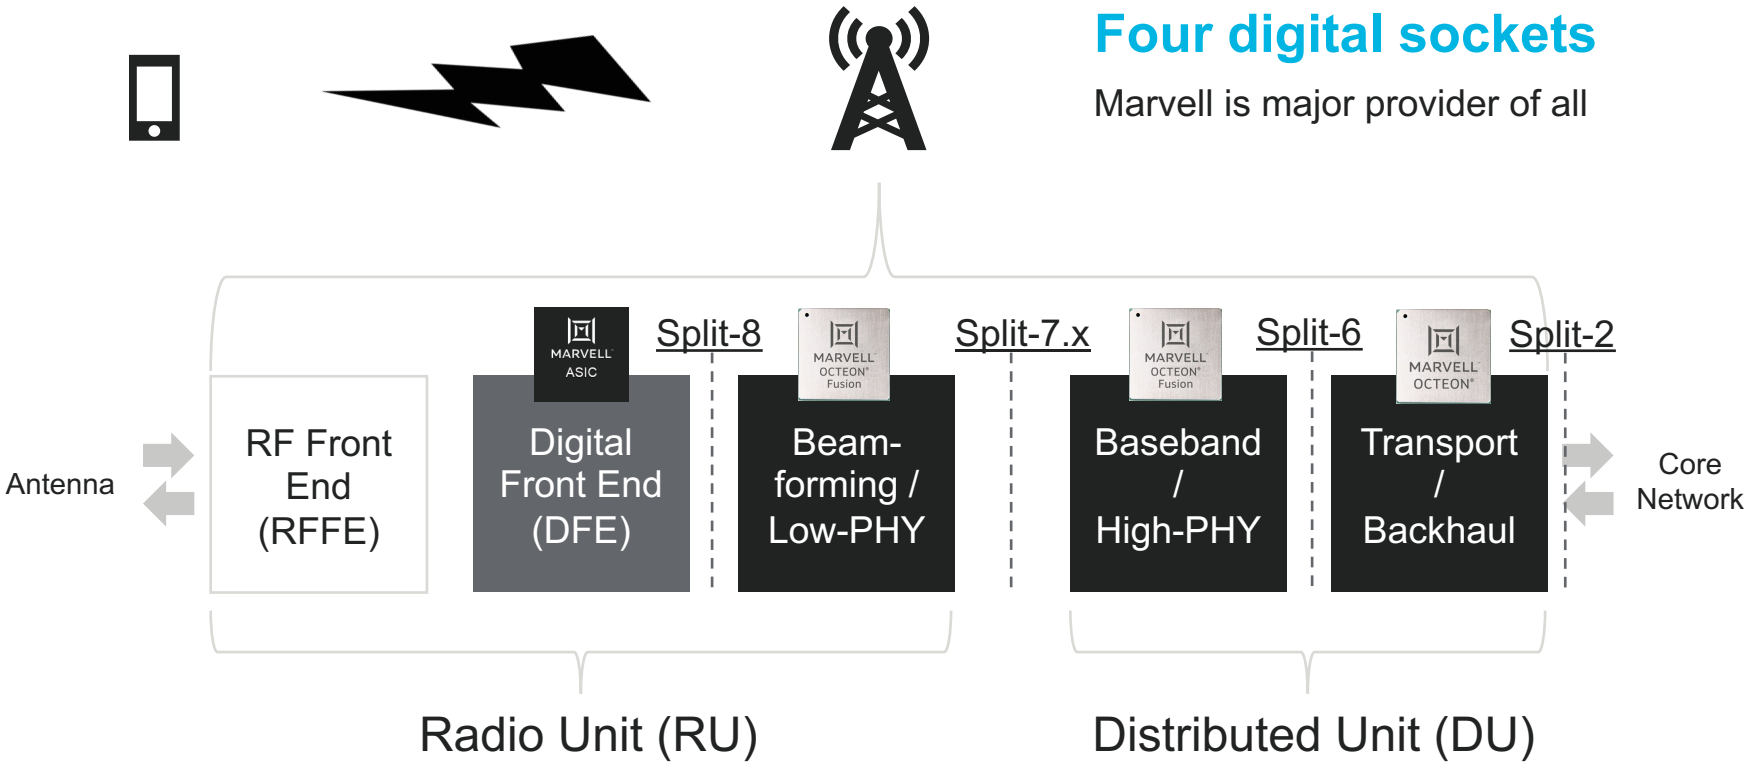

5G O-RAN solution: open, scalable, best-in-class

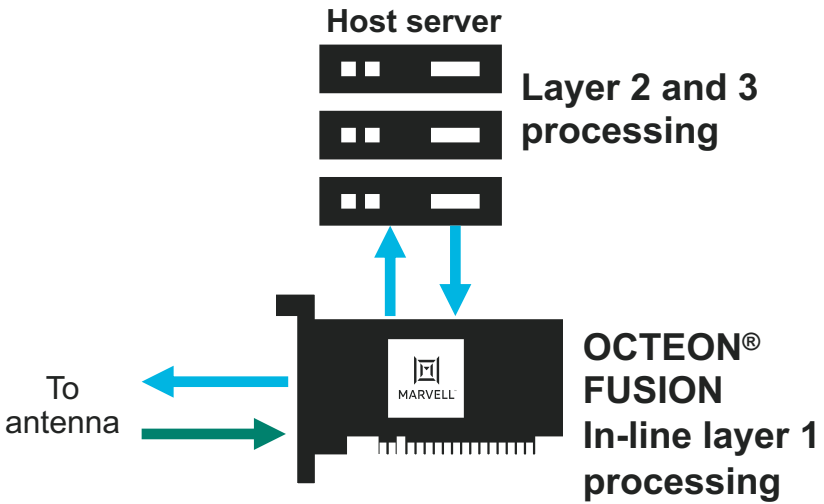

news release

vodafone

25 October 2022

VODAFONE AND SAMSUNG COOPERATE WITH MARVELL TO ACCELERATE OPEN RAN PERFORMANCE AND ADOPTION

news release

vodafone

25 October 2022

VODAFONE AND NOKIA PARTNER TO ADVANCE OPEN RAN ECOSYSTEM IN EUROPE

Vodafone and Nokia have agreed to jointly work on a fully compliant Open Radio Access Network (RAN) solution, marking a significant milestone for the mobile industry and a major boost to Europe's competitiveness.

The combination of Nokia's ReefShark advanced System on Chip (SoC) technology, developed in cooperation with Marvell, with standard Commercial-Off-the-Shelf (COTS) servers will enable the Open RAN system to reach functionality and performance parity with traditional mobile radio networks. Nokia's ReefShark SoC boosts the Layer-1 processing capability, which is necessary to connect many users to the mobile base station and support high levels of mobile data traffic.

Data Center use cases

1 Network offload

Accelerate networking functions

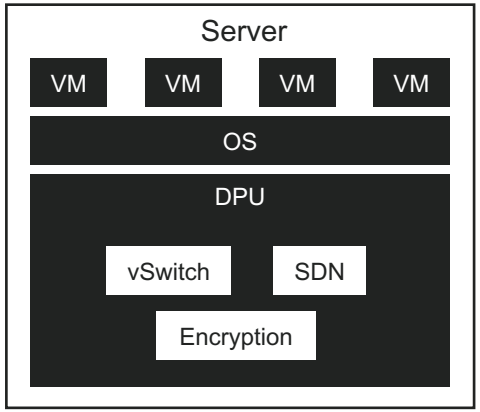

2 Storage offload

Accelerate storage functions

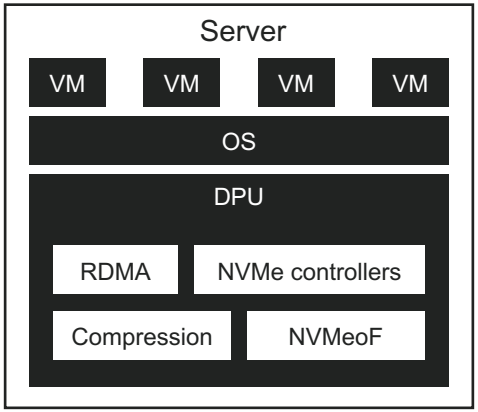

3 Security offload

Isolate tenants from host

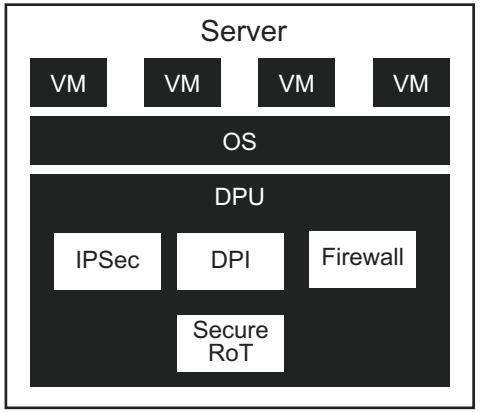

DPU delivers performance, programmability & lower TCO

Enterprise use cases

SMB router, SDWAN, gateway

Control and data plane

High-end router, Firewall/security data plane

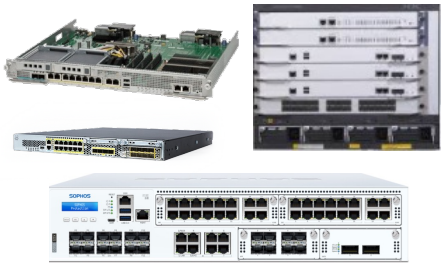

Data plane only

Switch, WLAN controller Line card/controller

Control plane only

Benchmarks

Compute performance

- SPECINT2006
- SPECINT2017
- CoreMark

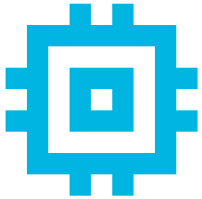

Memory subsystem

- LMbench
- Stream

Packet processing

- TestPMD
- L3 Forwarding
- iPerf/Netperf

Application benchmark

- IPSec gateway
- kTLS
- Open offload
- SNORT/Hyperscan
- NVMEoF
- ML Inference

200Gbps open offload performance

- Full Firewall VM running on Host server
- Marvell DPU as fast path engine - maintain route/firewall cache based on OpenOffload

CN98xx 2Ghz, 2x100G						
Data plane cores	64B		512B		1518B	
	Mpps	Gbps	Mpps	Gbps	Mpps	Gbps
2	3.54	2.38	3.54	15.1	3.46	42.5
10	15.7	10.6	15.7	67.1	15.6	191.9
30	43.1	28.9	43.1	183.5	16.2	200

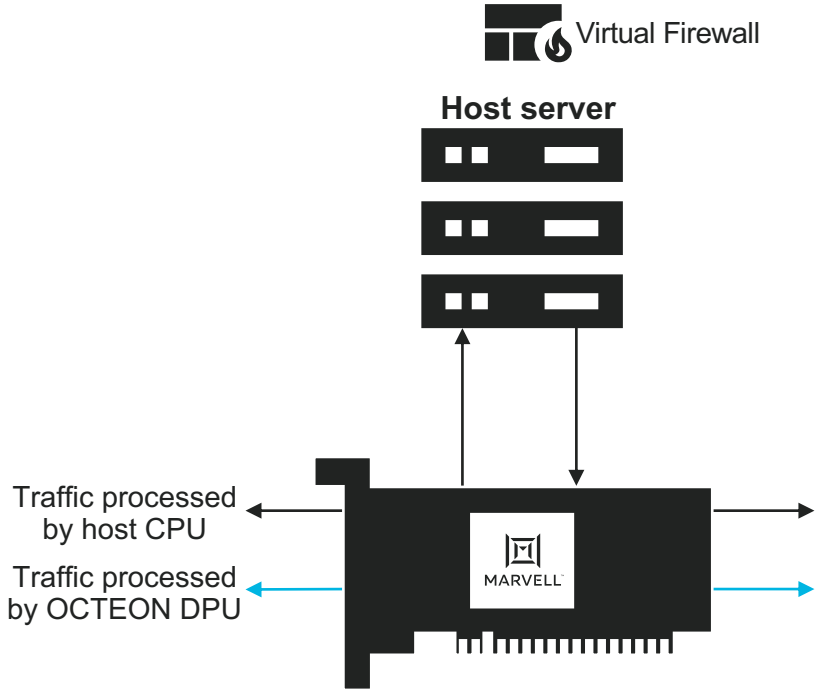

Integrated ML engine

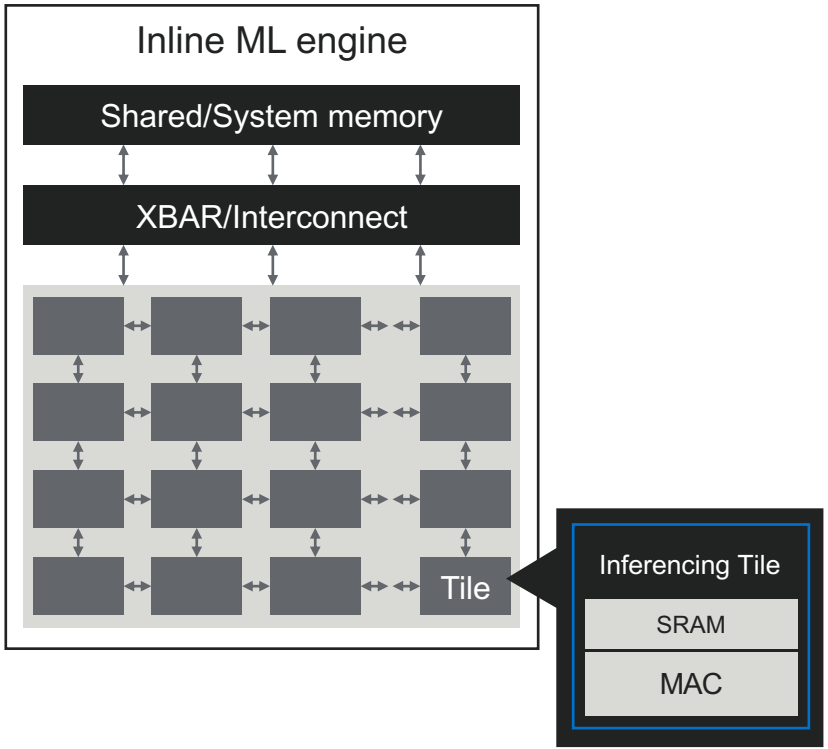

- **Best-in-class DPU inferencing**
  - Directly in the data pipeline
  - Each ML tile contains private SRAM
  - Ultra low power
- **Up to 100x performance vs SW**
  - Supports Int8, FP16
  - Accelerated Tanh and Sigmoid activation functions
- **Use cases**
  - Threat detection
  - Context-aware service delivery
  - QoS
  - Beamforming optimization
  - Predictive maintenance

Summary

1. Most widely deployed DPU, shipping since 2005 for data center, enterprise and carrier use cases

2. Software-defined infrastructure requires hardware acceleration – OCTEON® portfolio has right accelerators to deliver best solution TCO

3. Unified software stack built on open source frameworks and benchmarks demonstrate leadership across broad workloads

MARVELL

Q&A

MARVELL

Thank You

Essential technology, done right™

Session 5: Design

- **Increasing Data-Center Reliability by Predicting Defects and Monitoring Aging**
  - Marc Hutner, Senior Director of Product Marketing, **proteanTecs**
- **A New Way of Measuring Heterogeneous-SoC Performance**
  - Hanan Moller, Technical Director, **Siemens**
- **Advanced Semiconductor Manufacturing Technologies for Quantum Computing**
  - Anthony Yu, Vice President, Silicon Photonics Product Management, GlobalFoundries

proteanTecs
Electronics visibility. From within.

Increasing Datacenter Reliability by Predicting Defects and Monitoring Aging

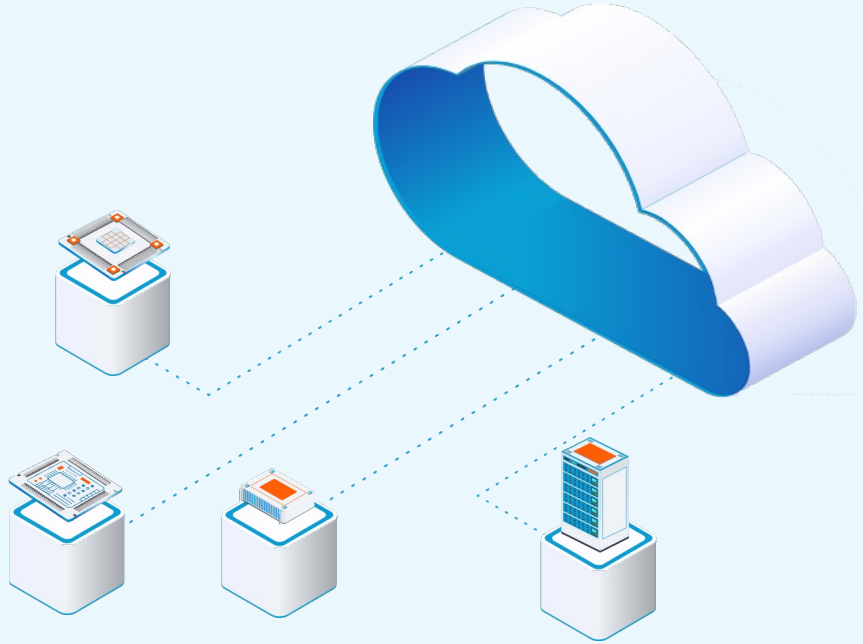

November 2022

Marc Hutner | Senior Director of Product Marketing
www.proteanTecs.com

The Linley Group
Fall Processor Conference 2022

proteanTecs at a Glance

Deep data health & performance monitoring for advanced electronics

Founded in 2017 by industry leaders and **co-founders of Mellanox**

Addressing industry-wide challenges of **scale**

50+ design wins. In-use & production-proven

Customers in multiple key segments including **Datacenter, Automotive, Communications, and Mobile**

New category with a **multi-disciplinary** approach

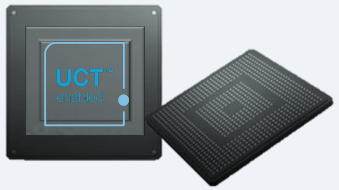

In use and production ready in advanced FinFET nodes from leading Foundries

Global Footprint

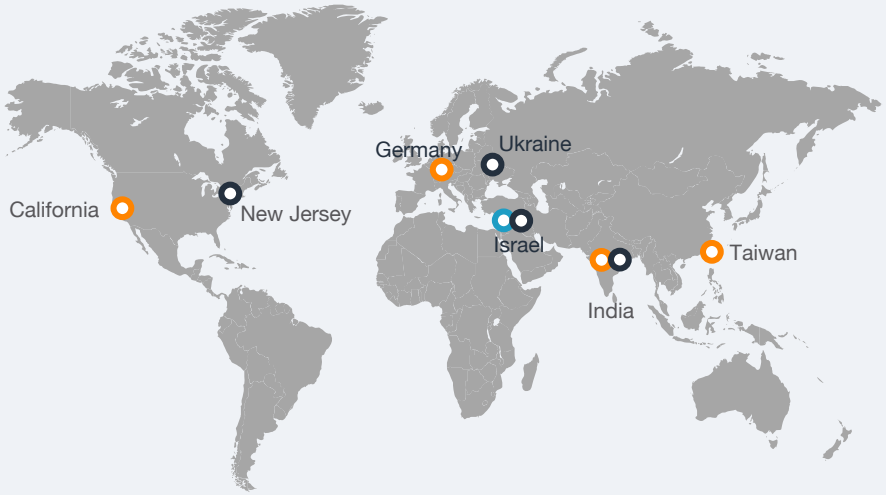

Alliance Members:

tsmc Open Innovation Platform®

SAMSUNG SAFE™

Increasing Challenges in Advanced Electronics

Increased thermal density
High performance applications with greater transistor density

Advanced Packaging
Micro-bump cracks and voids, silicon warpage and bridge-shorts on substrate

Stronger impact of degradation & failure mechanisms
In-mission manifestation of physics-of-failure becomes common

Latent defects
Escapes of marginal or latent defects due to aggressive design architectures and advanced process technologies

SW impact on HW
AI/ML driven applications with dynamic mission profile create ongoing risks

Real Examples of Datacenter Failures

The New York Times

"Increasing complexity in processor design, smaller transistors, 3D chips and new designs that create errors only in certain cases all contributed to the problem"

"Several years ago, Google production teams began \to report errors that were maddeningly difficult to diagnose. Calculation errors would happen intermittently"

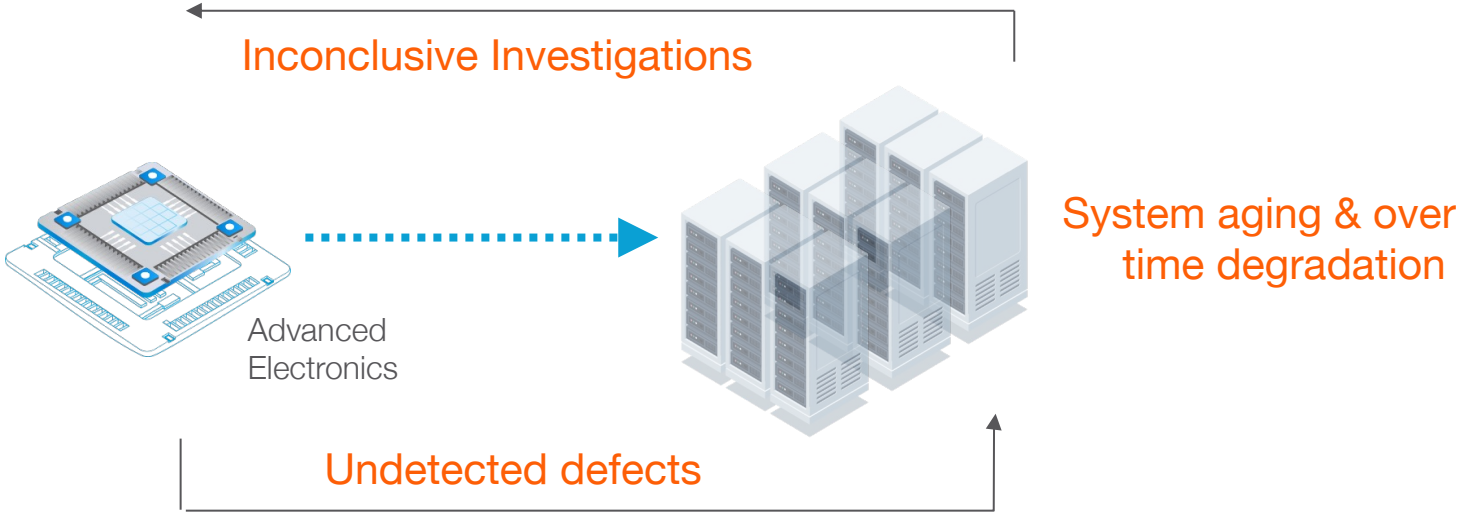

Market leaders are looking for answers from the industry

facebook.

"Silent data corruption due to silicon latent defects and aging. (1,000 DPPM!)"

Microsoft

"Frequent unexplained HW failures with "No Issue Found" at high rates"

Google

"Rare, short-term computational errors on systems that passed all manufacturing tests successfully"

1. "Silent Data Corruptions at Scale"; Harish Dattatraya Dixit Et. Al., Facebook, arxiv.org/abs/2102.11245, Feb. 2021
2. "Cores that Don't Count"; Peter H. Hochschild Et. Al., Google Research, Jun. 2021
3. "Improving Cloud Scale Hardware Fault Diagnostics", Neeraj Ladkani, Rama Bhimanadhuni, Microsoft, Mar. 2020 OCP Virtual Summit
4. The New York Times, "Tiny Chips, Big Headaches", February 16, 2022

The Chip as a Smart Datacenter Monitor

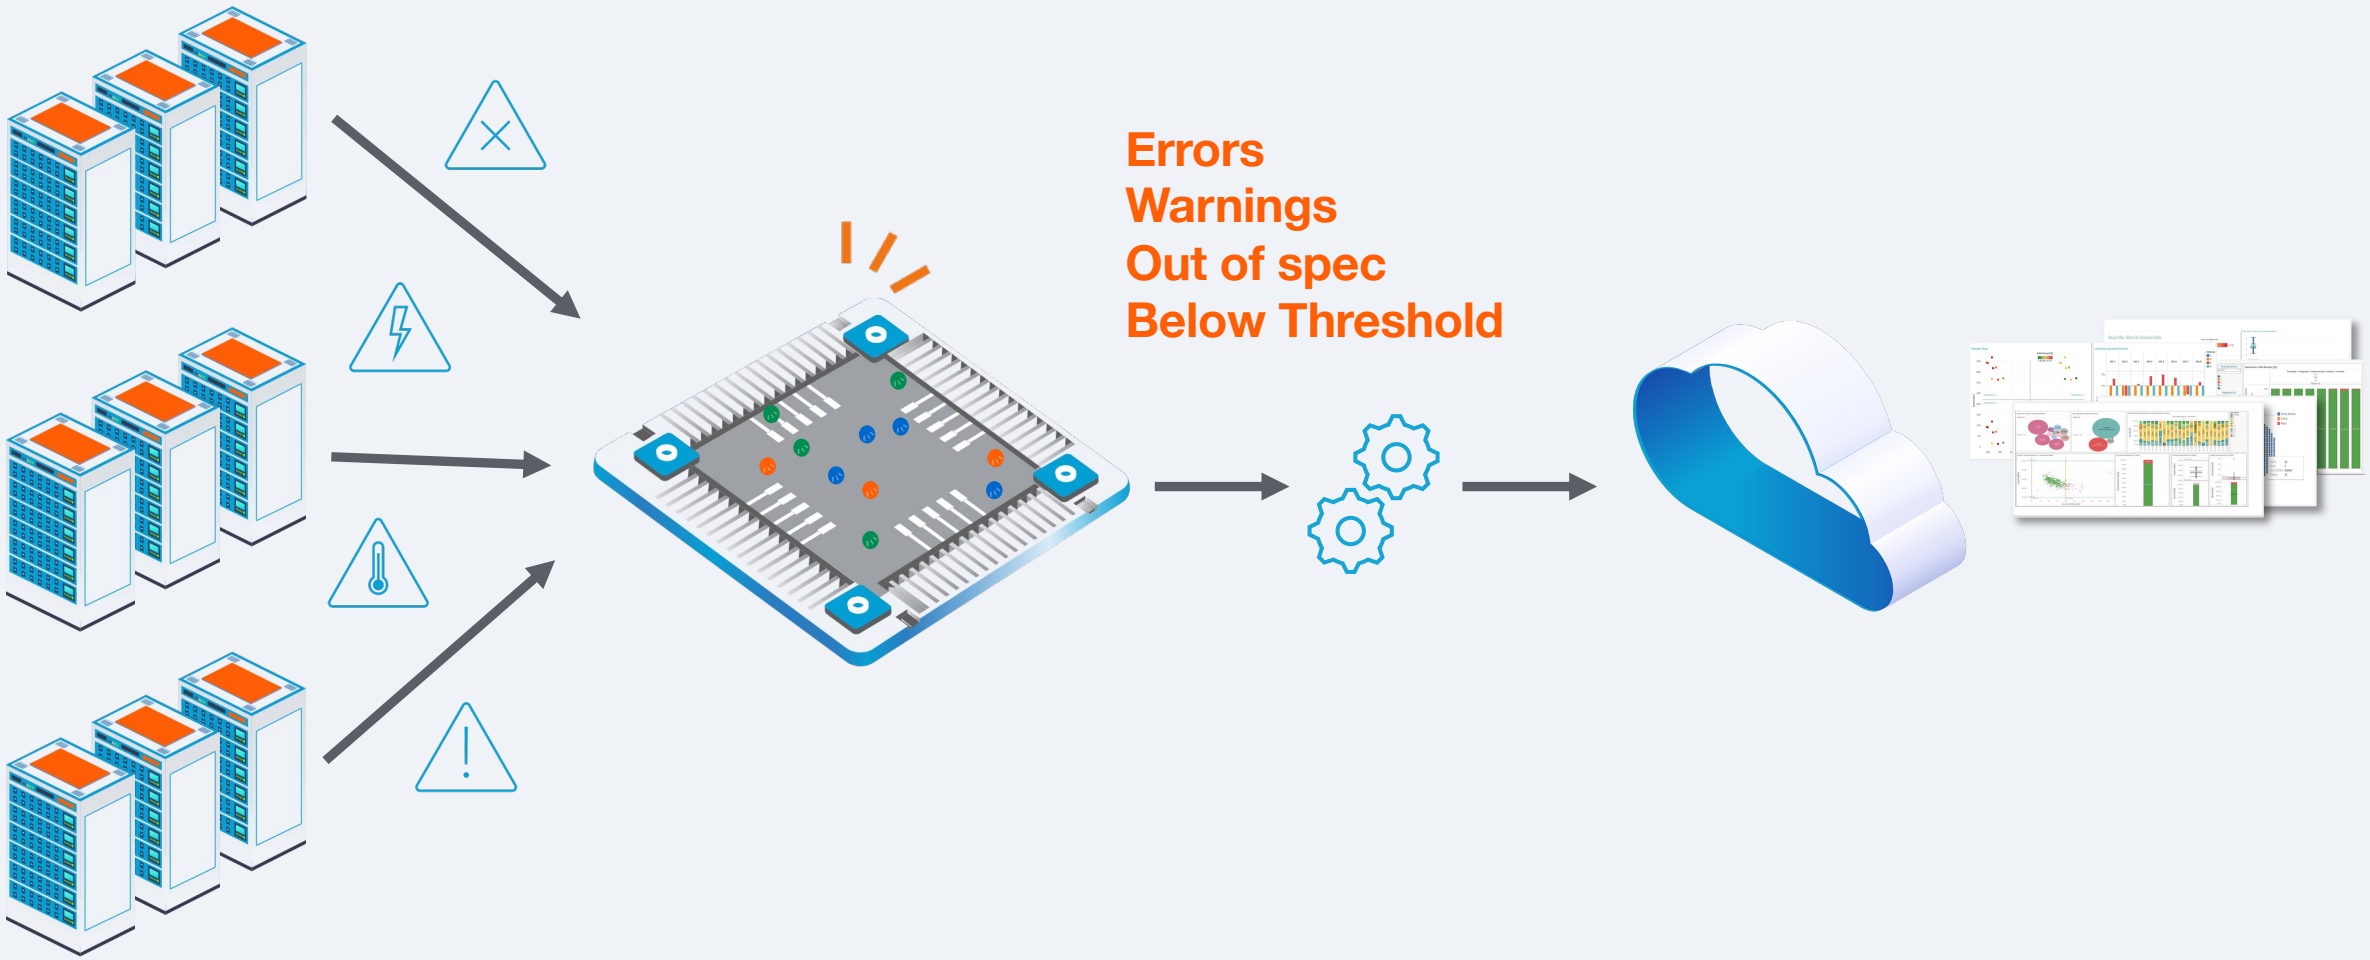

Driving Visibility and Actions at Every Stage

Chip Production

- Higher confidence volume ramp
- DPPM reduction
- Performance yield
- Power reduction

System Production

- Performance optimization
- Environmental monitoring
- Quality monitoring
- SW<>HW optimization

Applications for In-Field Monitoring

- Performance monitoring
- Alerts on faults before failure
- Power/performance/health strategies

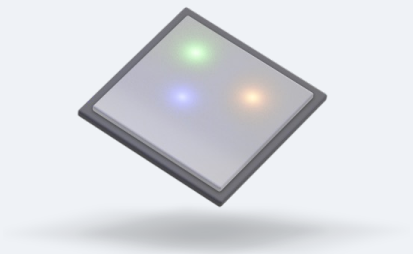

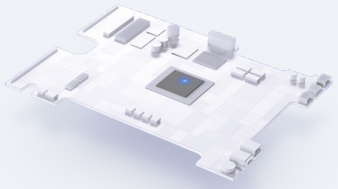

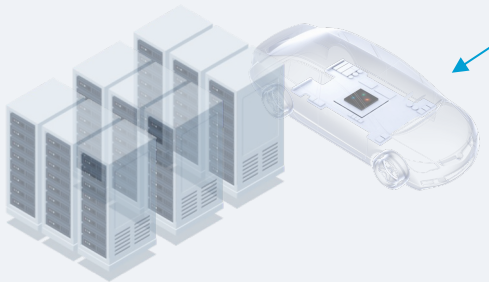

Machine Learning & Advanced Analytics

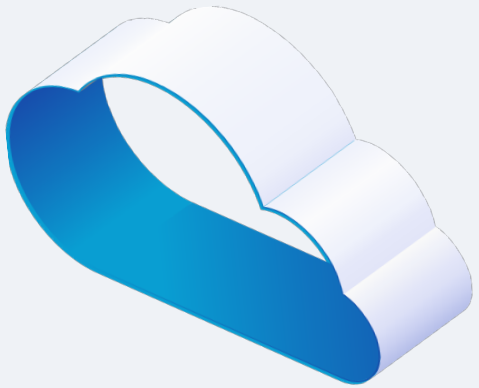

Powered by deep data

10X Lower DPPM

Faster Time to Market

Reduced Costs

Failure Prevention

Universal Chip Telemetry™ (UCT)

On-chip Agents built for analytics

- Parametric measurements
- High coverage & high resolution
- Minimal PPA impact
- Operate in mission-mode
- Sense the surrounding electronics
- Application optimization to HW

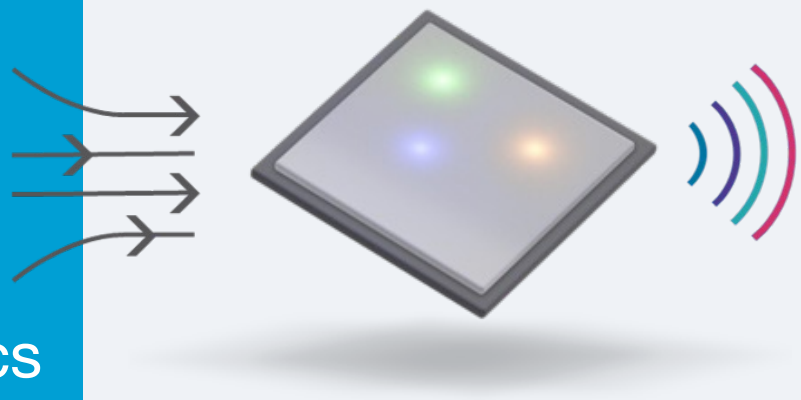

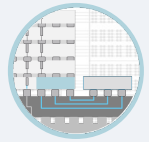

Interconnect Performance Monitoring

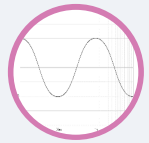

Operational Monitoring

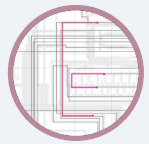

Performance and Degradation Monitoring

Classification and Profiling

Realtime monitors

Automated Health Monitoring for Production

Automated Health-Monitoring Production Workflow

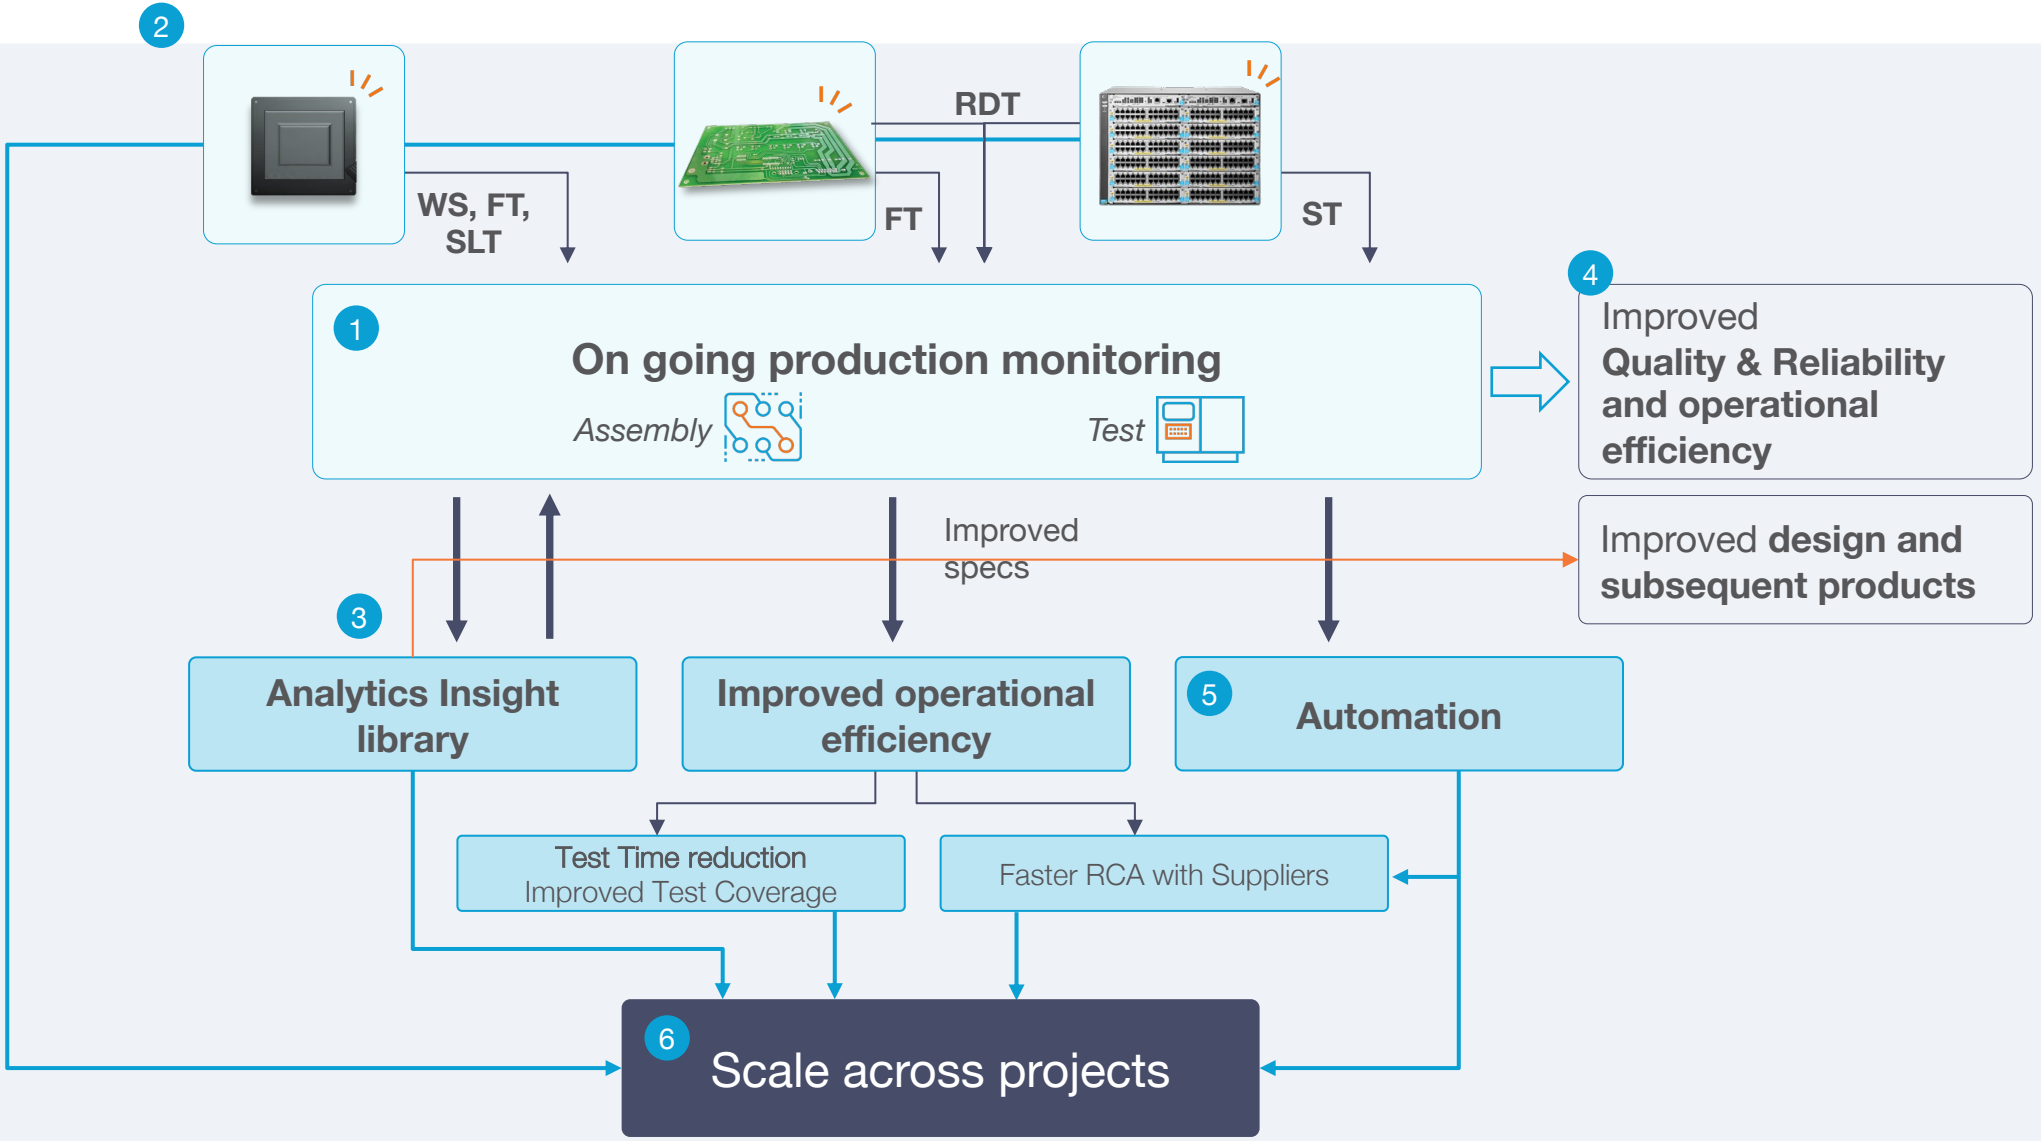

Use cases for Chip Production

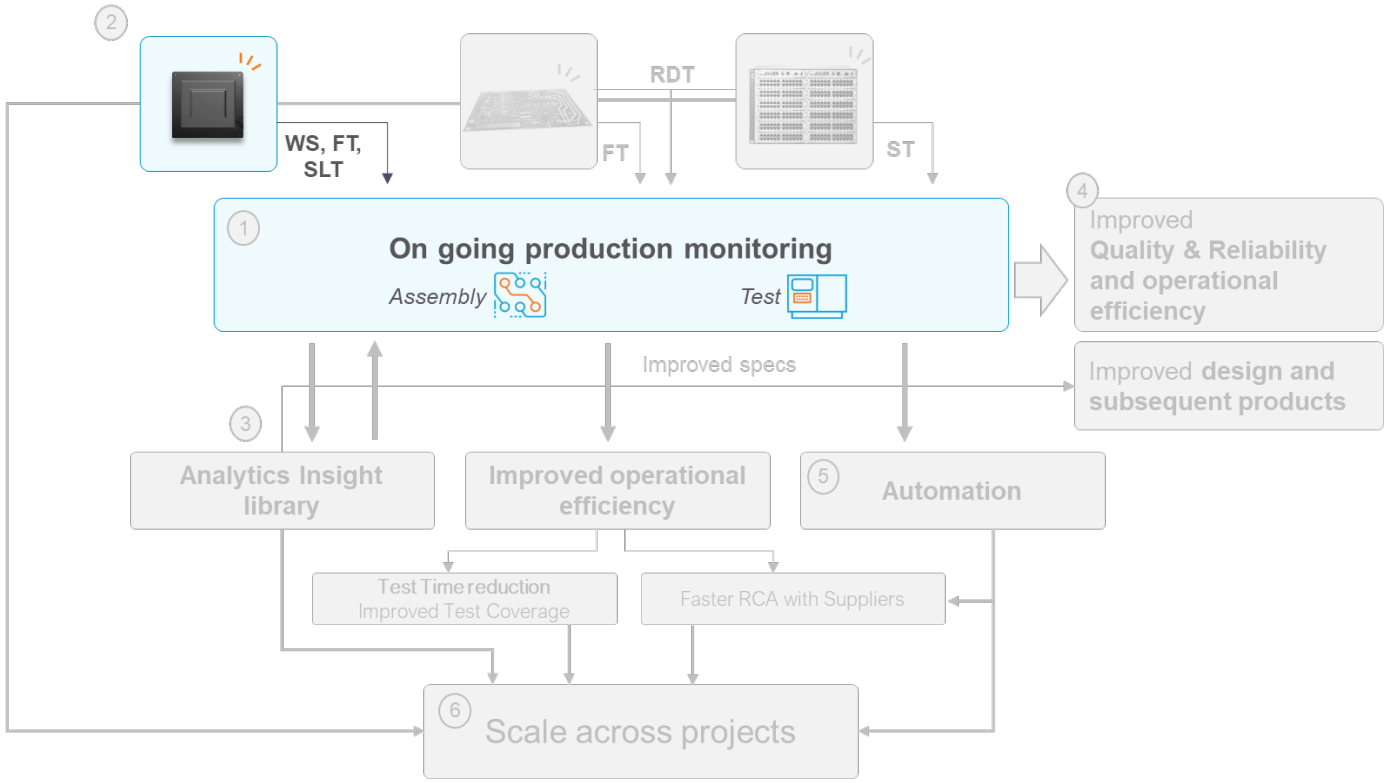

Personalized Device Assessment

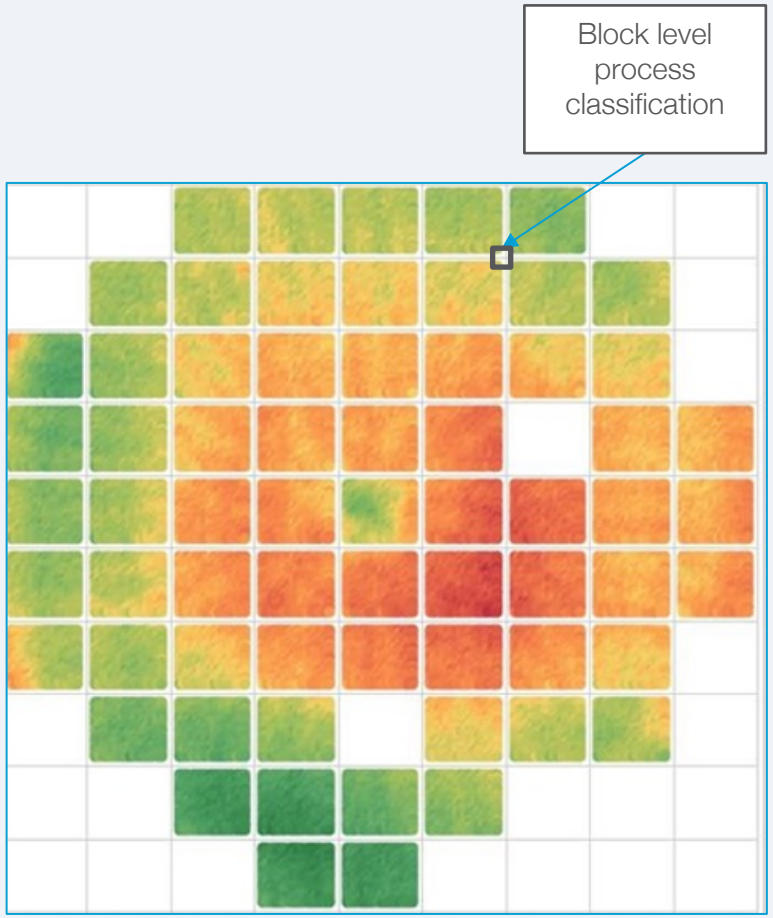

High granularity process parameter measurements (wafer, die, block)

From a pass/fail per-chip view to a **parametric per-block view**

1. Post-to-pre silicon correlation
2. Inferred process parameters
3. High resolution OCV
4. Vmin/ Fmax predictions
5. Smart material selection
6. Performance yield optimization
7. Bin planning
8. Advanced debug

Prediction Based Per-Die Outlier Detection

- Personalized chip assessment
- Multi dimensional outlier detection
- Yield reclamation
- Fast RMA with pinpoint Root Cause Analysis

BKM: Population based

proteanTecs: Dynamic limit per chip

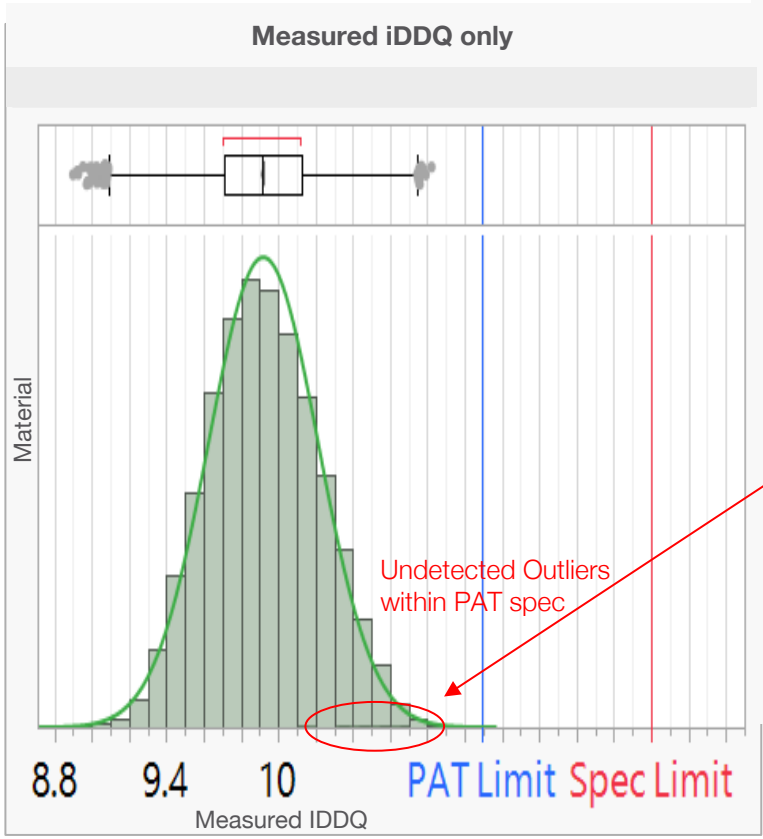

One dimension only identifies outliers ***outside*** of PAT spec limits (circled in red – unidentified outliers)

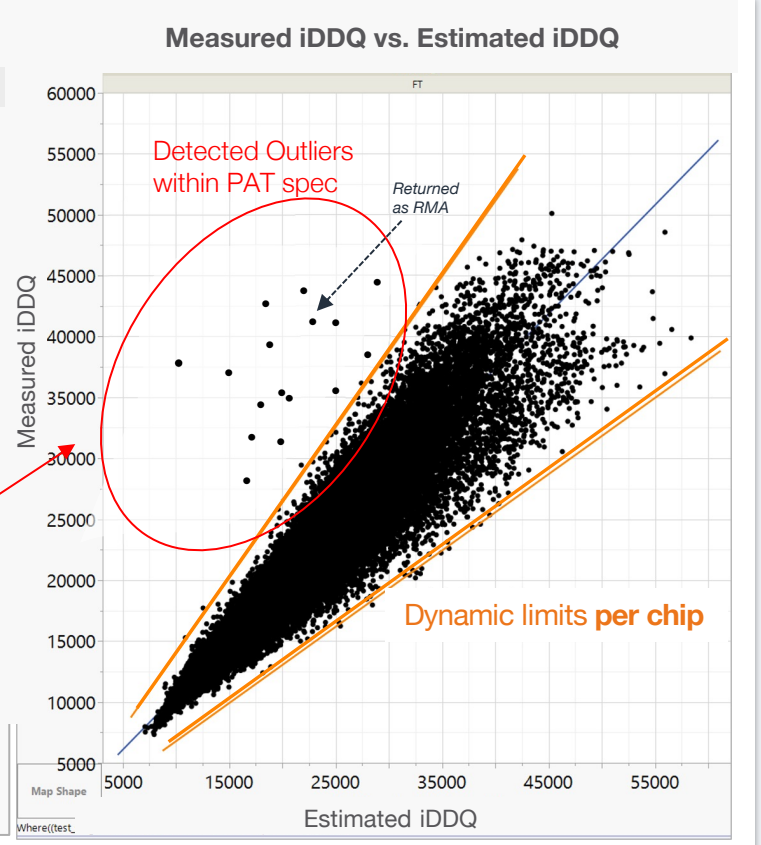

Additional dimension to identify outliers ***within*** PAT spec limits (circled in red – identified outliers)

Margin Based Latent Defect Detection

- Detection of **extrinsic failures** (latent defects) to prevent RMAs
- Margin quantification of the **path delay of millions of paths** (X% coverage) provides **parametric measurements**
- Sense and determine if the path delay is different than expected and mark the unit as **suspected/ marginal**
- Could be caused by defects in frontend processes (transistor defects) or backend processes (missing vias, small bridging defects)

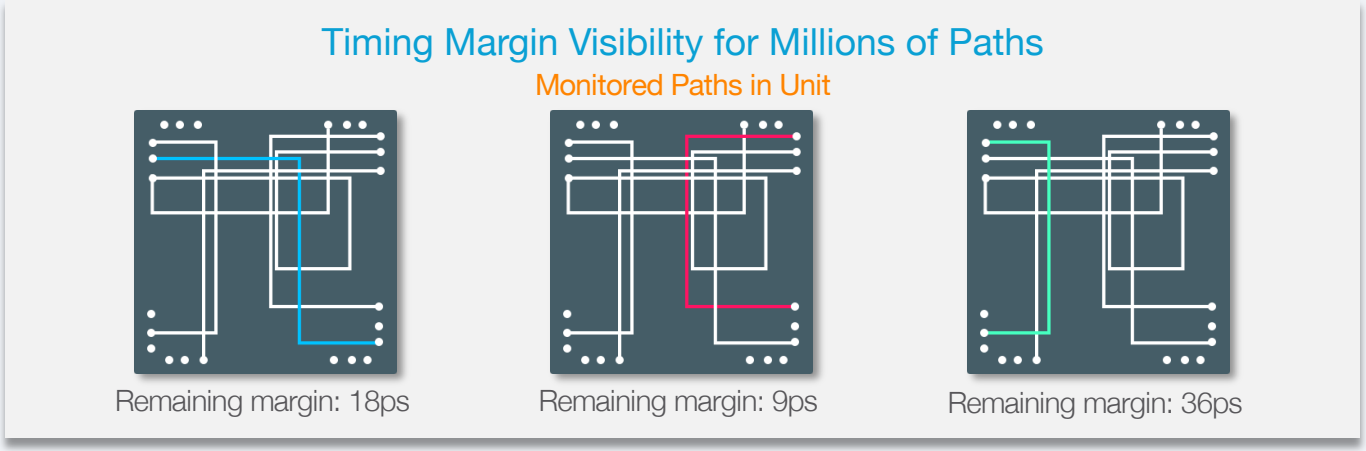

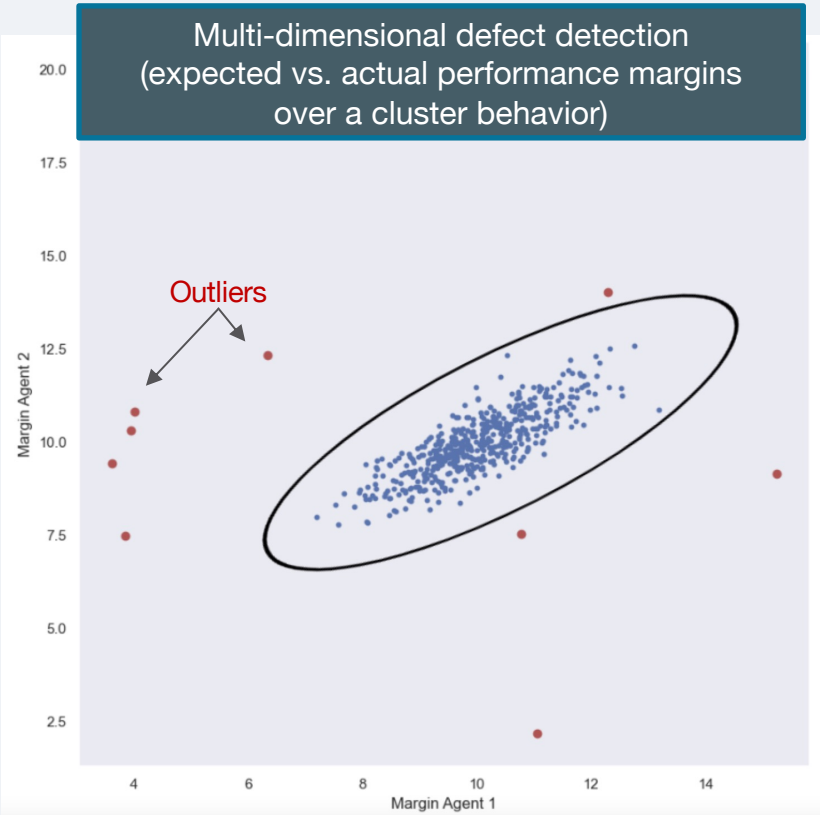

Use cases for System Production

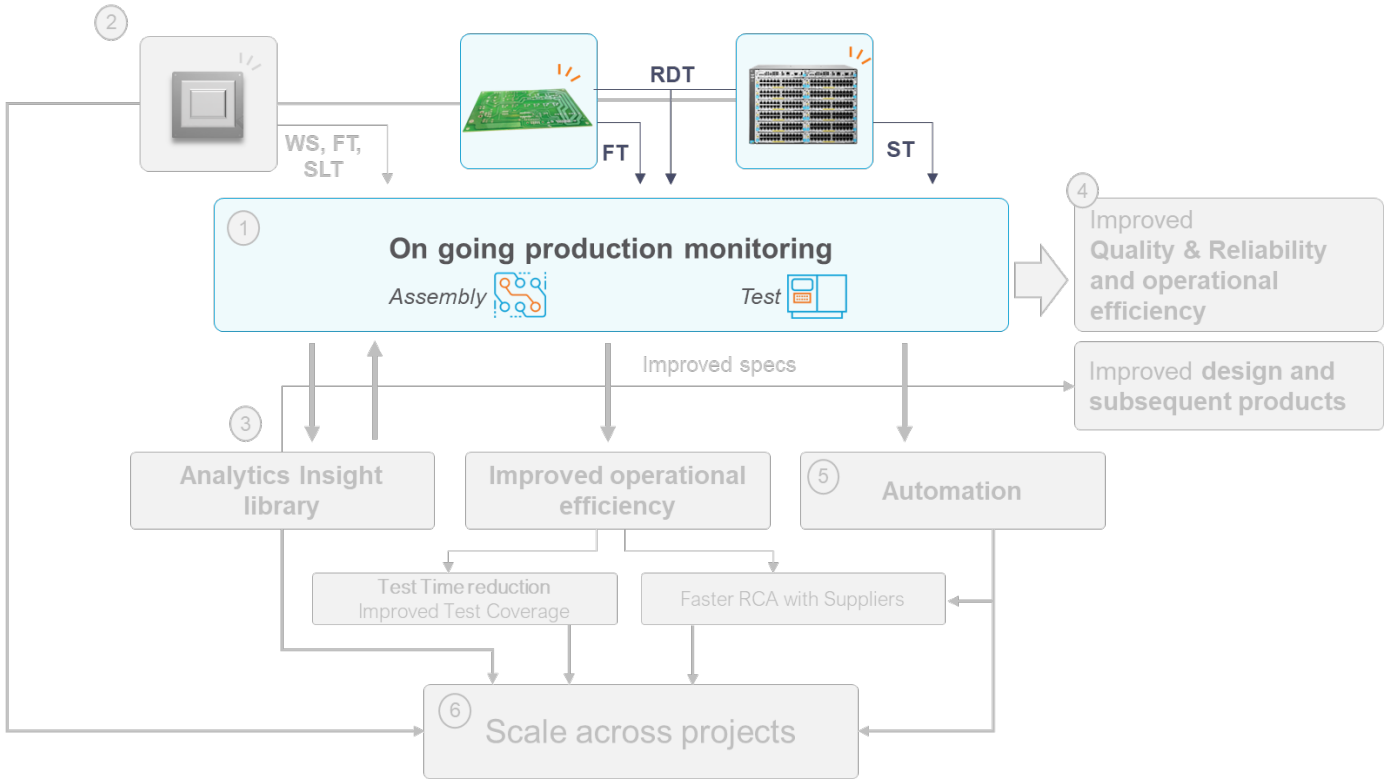

High Volume System Production

Typical board assembly line

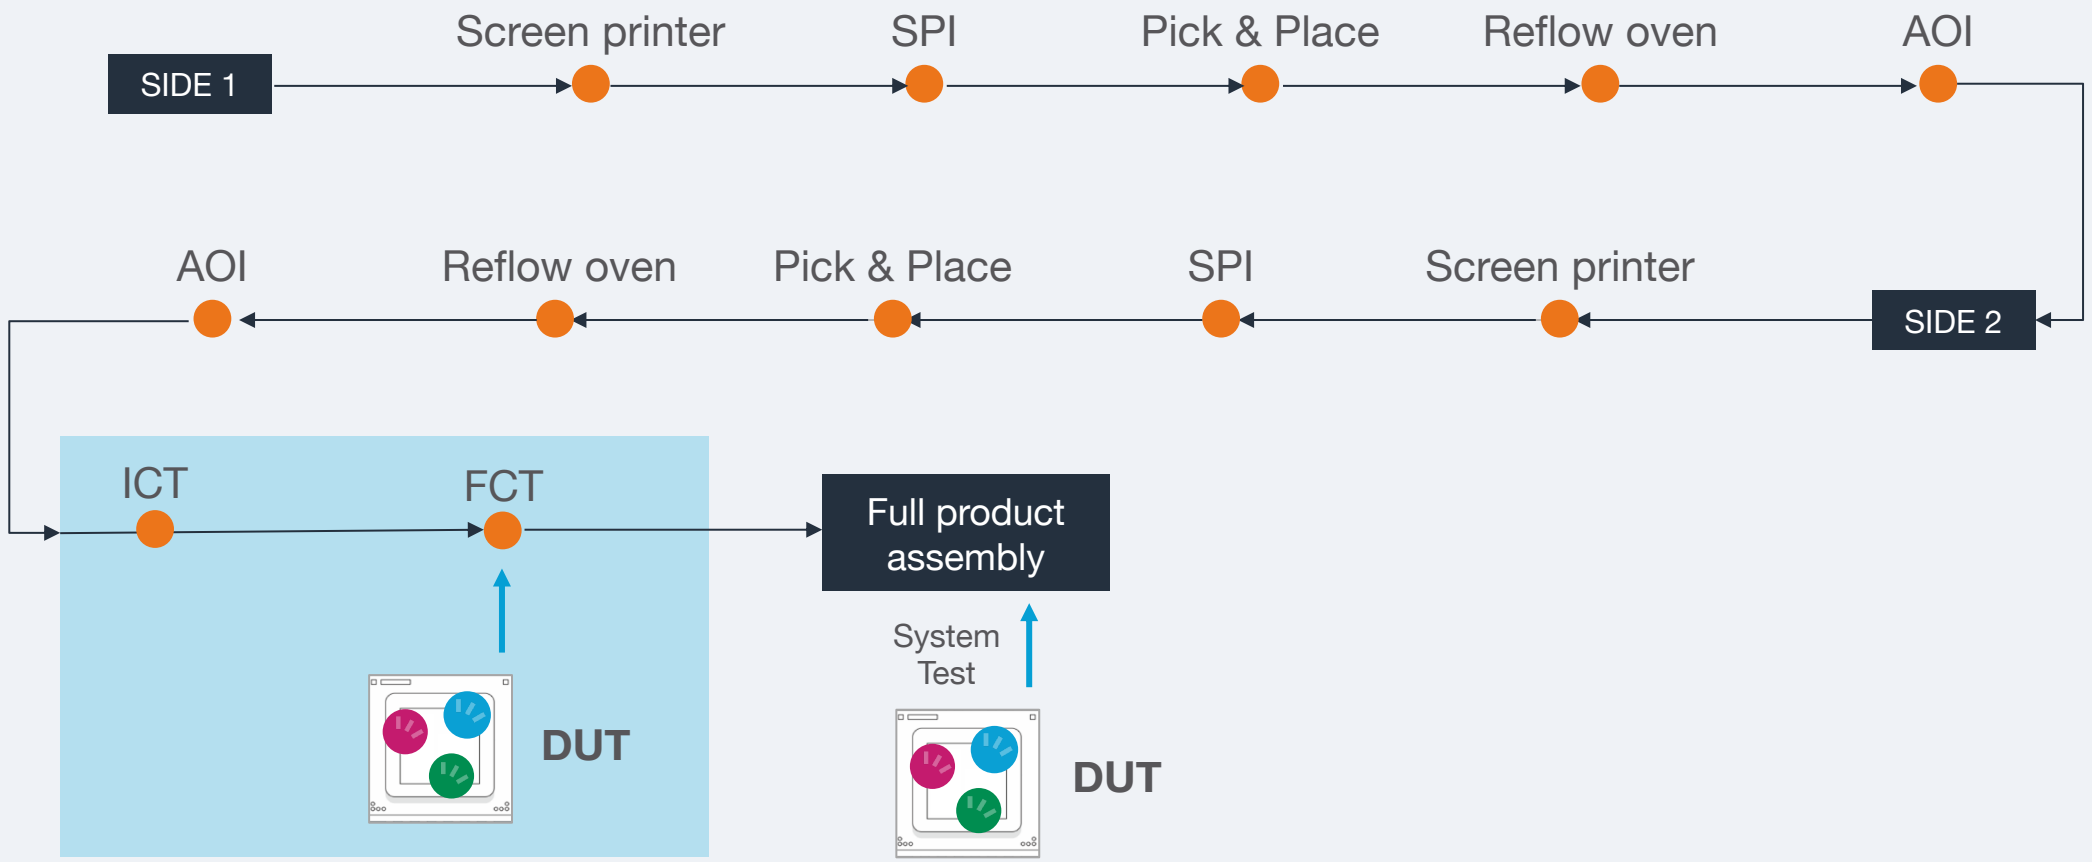

Fast Identification of Software Performance Issue

- Example: Lower margin detected at System Level: root caused to the effect of the application functional tests as measured by the system level, with correlation to workload and operational measurements
- Demonstrates the impact of different system level applications on the chip performance

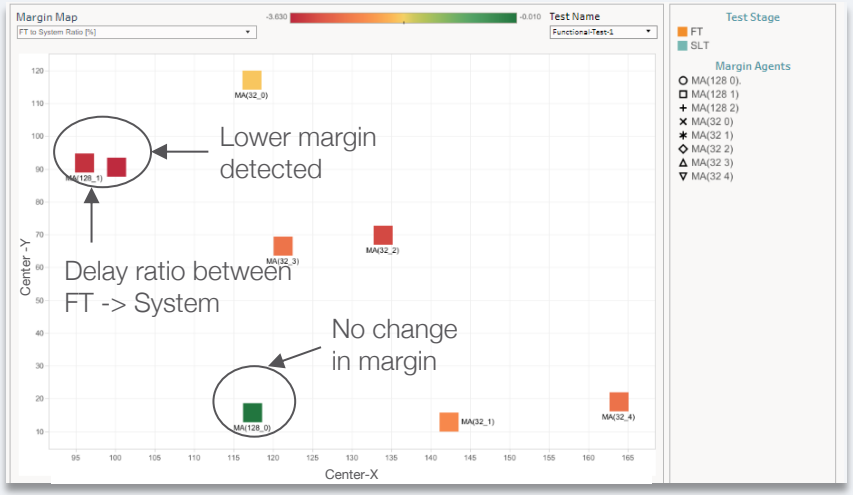

IC spatial view of Margin measurements

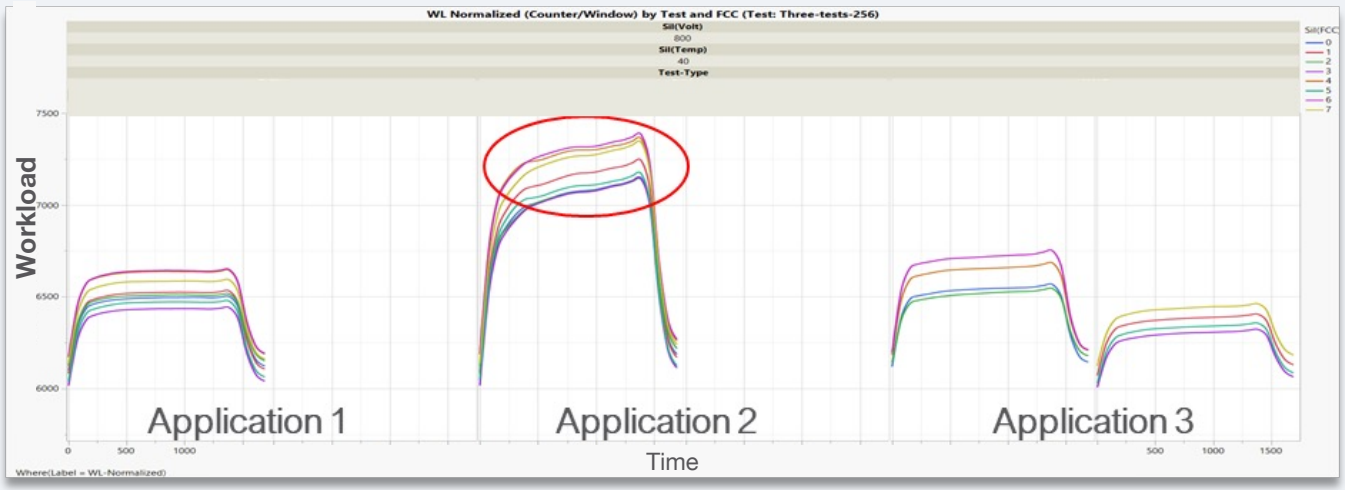

Workload root cause analysis

System Performance Monitoring

Agent data recorded throughout multiple phases of system operation:

- Configure
- Operation Start 1
- Operation End 1
- Idle 1
- Idle 2
- Operation Start 2
- Operation End 2
- Idle 3
- Idle 4

- Correlating and merging different Agent measurements
- Margins of millions of paths
- IR drop
- Workload/stress: V*T*Toggle rate
- Cycle to cycle clock jitter

One Margin Agent

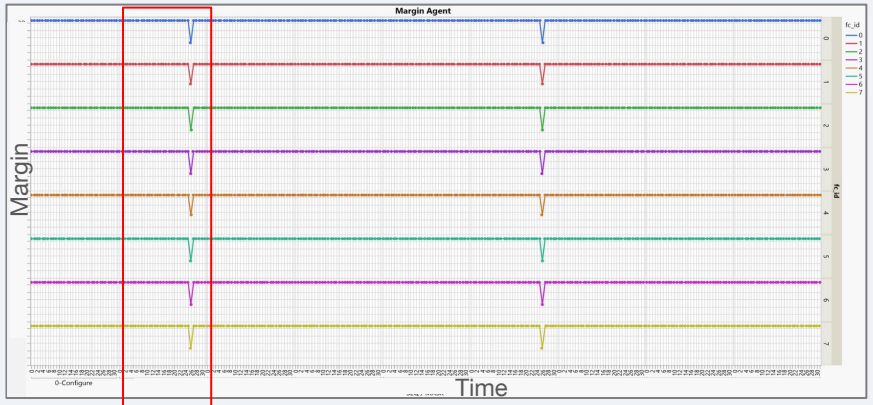

Workload Agent

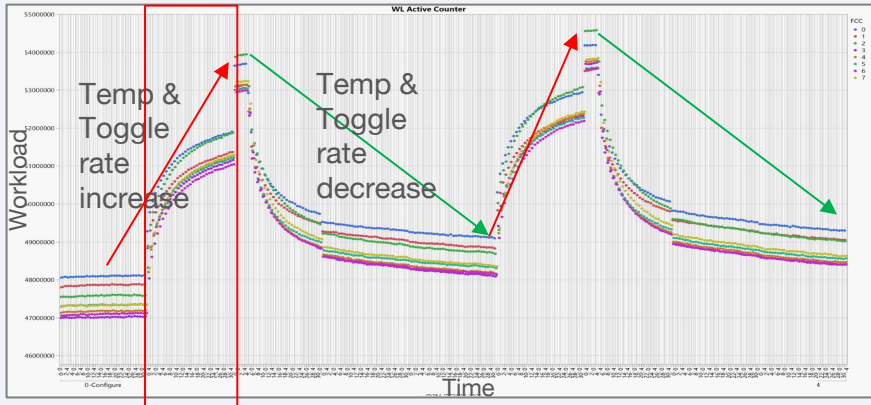

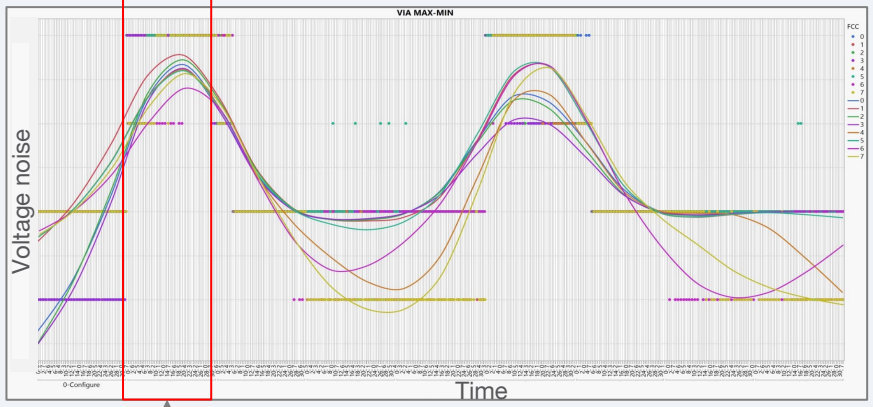

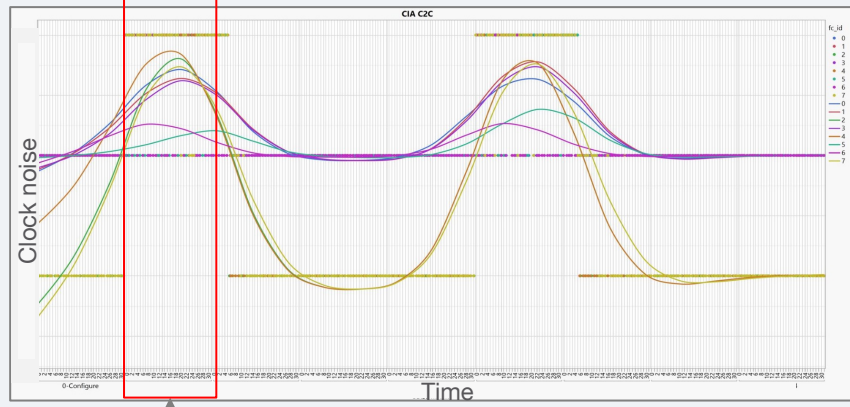

Voltage Integrity Agent

Clock Integrity Agent

Operation Start 1 Phase

Bring up & Production EOL software optimization and performance tuning

Correlation Across Stages/ Setups/ Tests

Accurate correlation & validation between stages: setup and test stress

Optimization of test conditions across all test stages (WS, FT, SLT, ST)

- Test limits and guard-bands, test equipment validation, etc.
- Accurate characterization of operational, environmental and stress effects

- ATPG structural test tuning - correlative to real application at system
  - Common data ‘language’ between ATE and system test

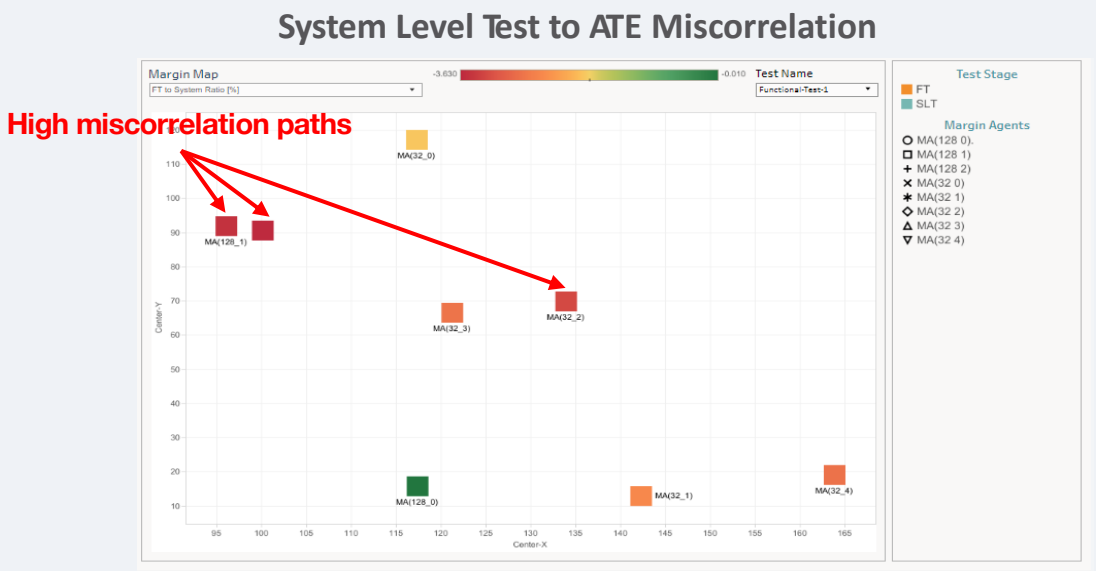

IC spatial view: Lower margin detected at SLT

RCA

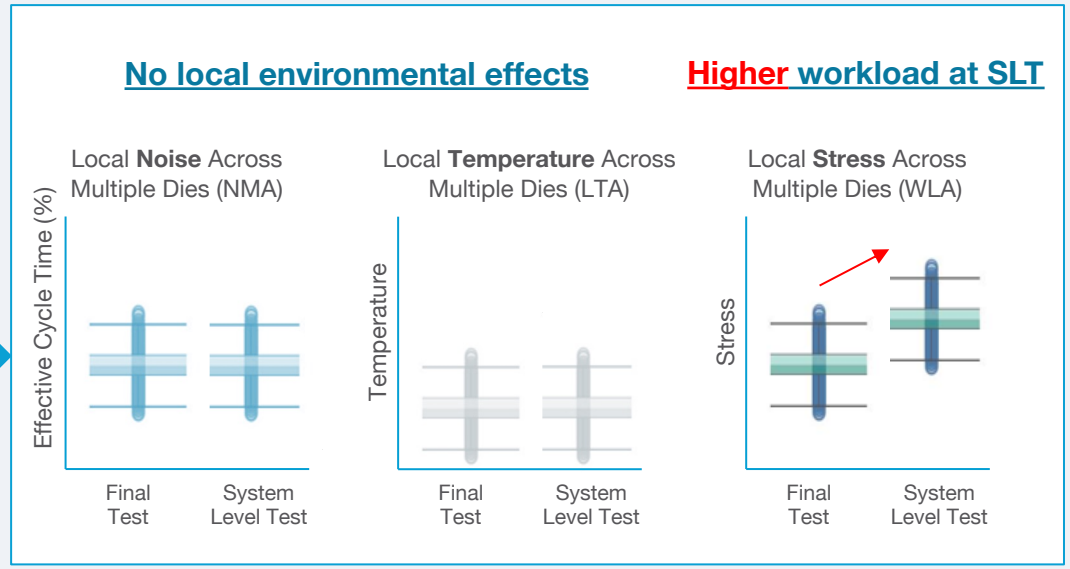

Example source of issues: Poor Performance Coverage at ATE

Automation at Test Production Line

- Edge libraries are deployed to production line test for real time decisions
- Insight monitoring and operational efficiencies (test time reduction, skip-test)

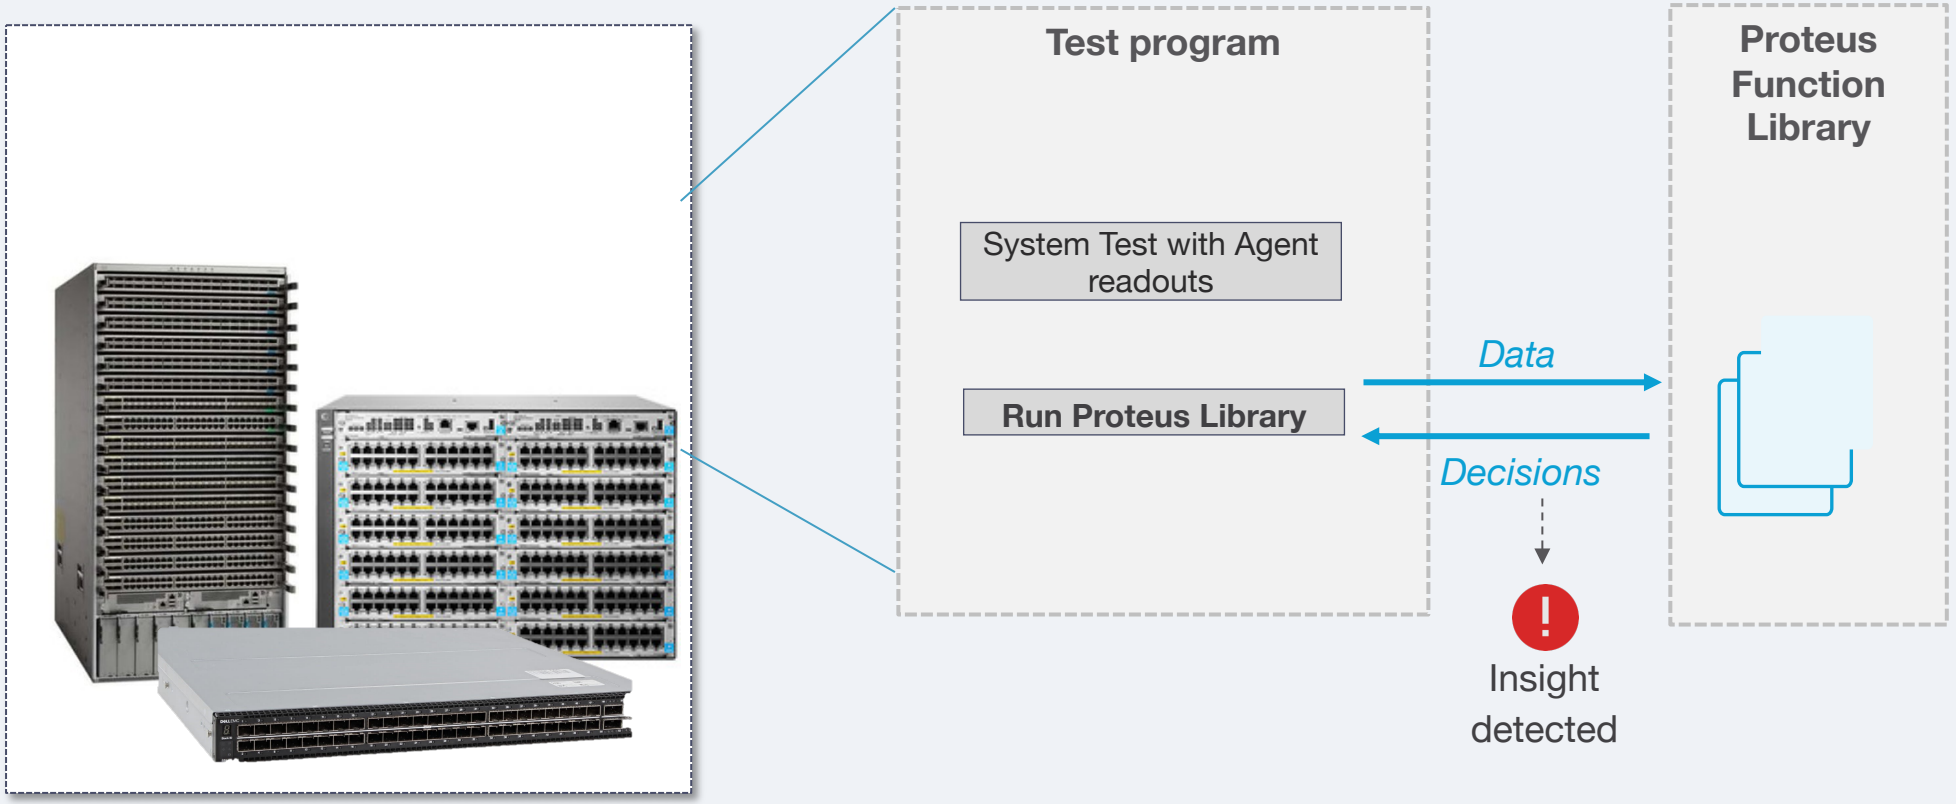

Health-Monitoring Automated Workflow for Production

New methodology for predictive analytics with ML-driven on-chip telemetry

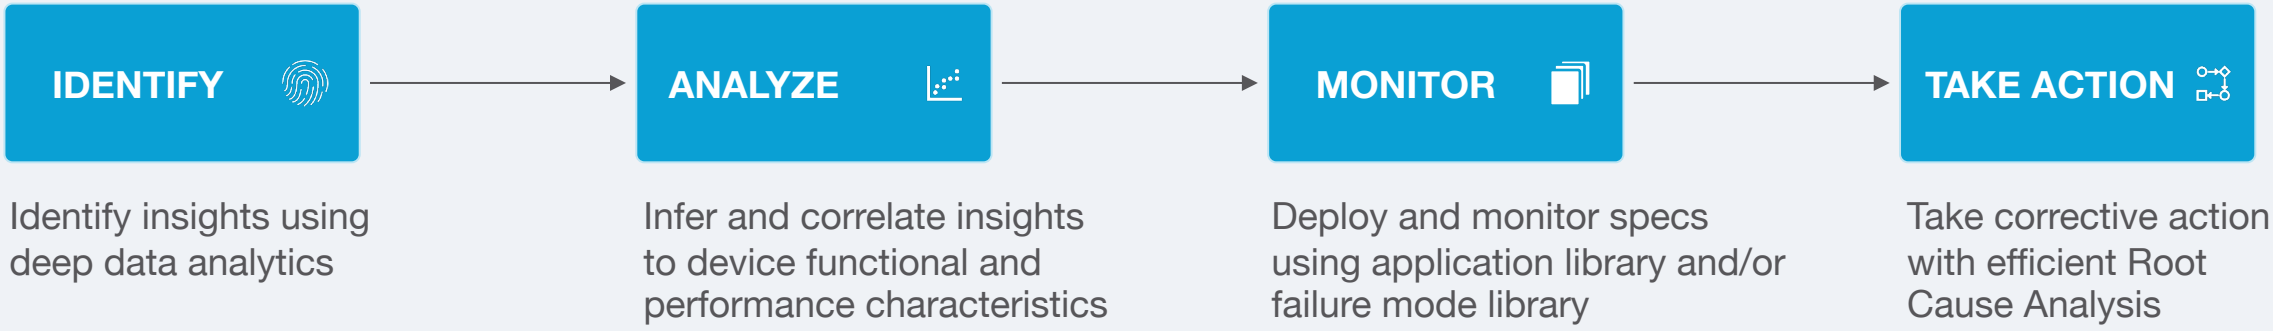

- **Structured monitoring with continuous learning loop for existing product**
- **Scale for subsequent products with automatic deployment of libraries and improved specs**

Use cases for In-Field Health and Performance Monitoring

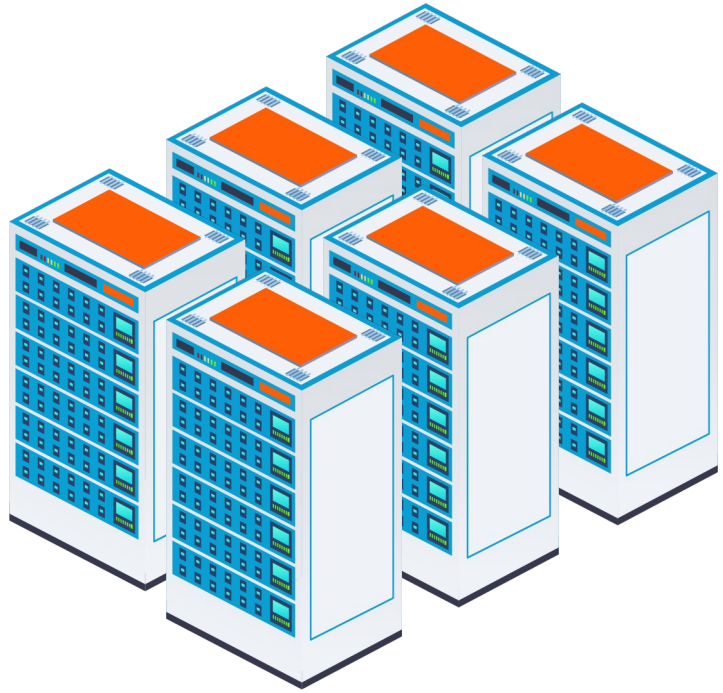

Continuous Performance Monitoring with On-Board Diagnostics

Edge application: Lowering the operational risks

- Misbehaviour → close to failure
- Manipulation → out of reference, security breach or tuning
- Misuse → out of spec, violation of operational requirements

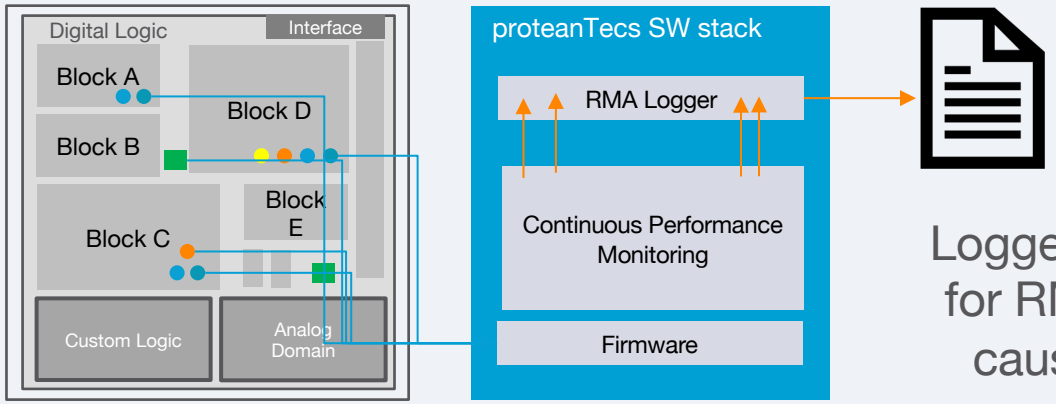

Logger application for RMA and root cause analysis

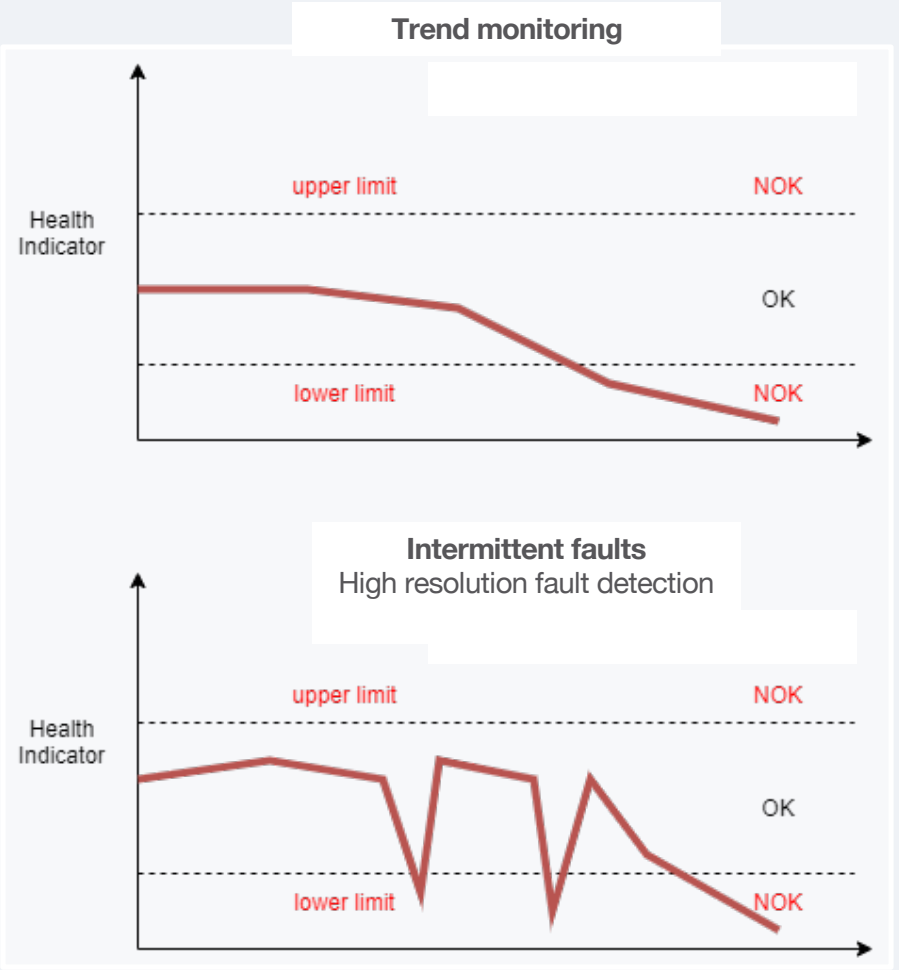

Performance Degradation Monitoring

Based on Margin and Workload- **Real operation**

- Margin measurements covering real paths
- Using frequency degradation as a precursor of failure (both **intrinsic and extrinsic** reliability)
- Covers latent defects (extrinsic reliability issues)
- **TTF prediction** based on Margin Agents measurements
- Works with Workload measurement to track usage changes

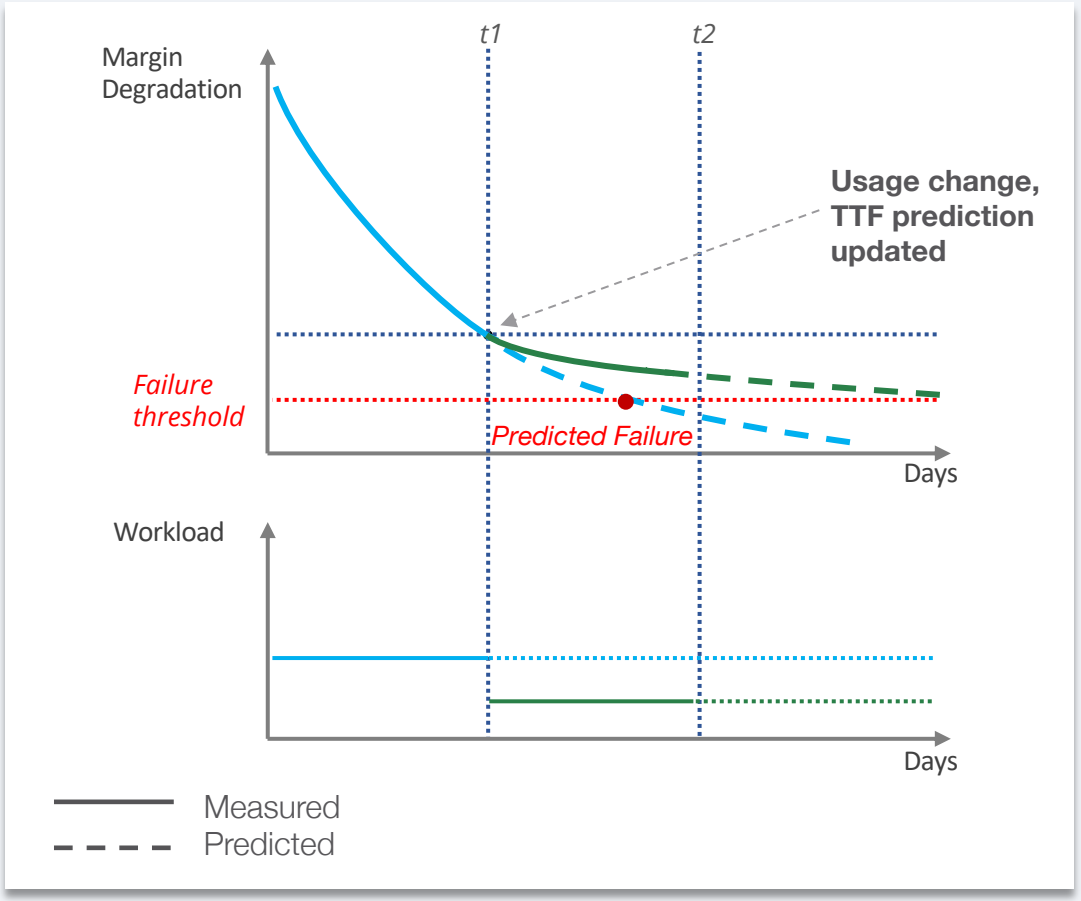

Continuous Performance Monitoring
Edge application with built in diagnostics

In-Field

- Timing margin failure
- Microcontroller overstress failure
- Local voltage failure (AC)
- Local – global voltage failure (AC)
- Power delivery network fault (DC)
- Power delivery network failure (AC low frequency)
- Effective local clock failure
- Clock source failure
- Thermal runaway
- Local thermal hotspot
- Connectivity failure

Different areas in the chip

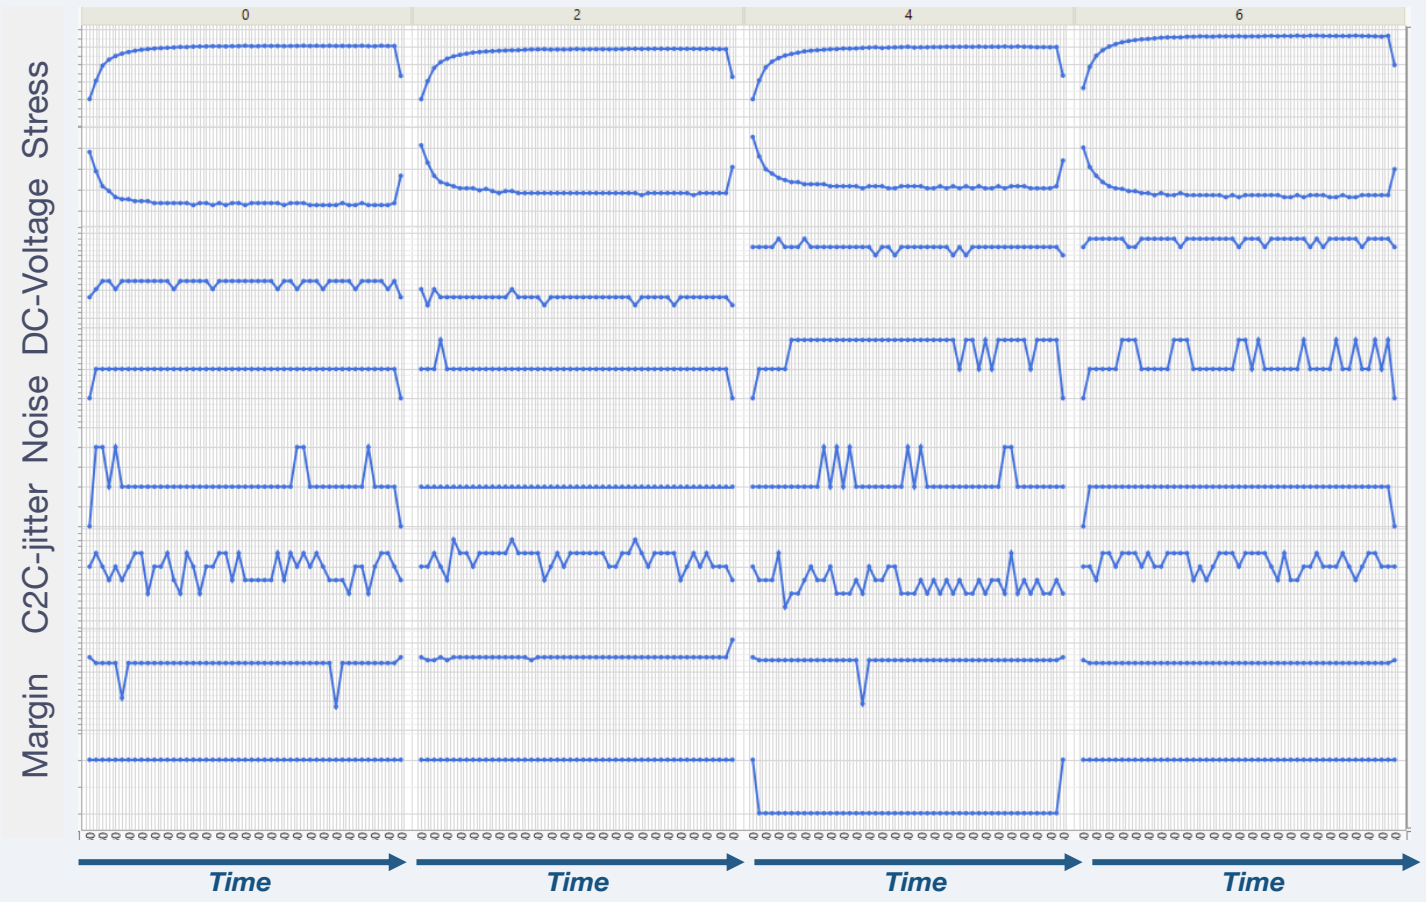

Optimize, Monitor and Scale

OPTIMIZE AND MONITOR

- Automated as part of production flow
- New RCA methodologies with suppliers
- New test and in-use monitoring methodologies
- Apples-to-apples cross stage correlation
- Derived specs improve subsequent designs
- Predictive performance monitoring

Quality Defect Detection
Faster Time to Market
RMA Reduction
Increased Service Availability

Reuse libraries for other Chip and System projects

SCALE

Thank you.

marc.hutner@proteanTecs.com

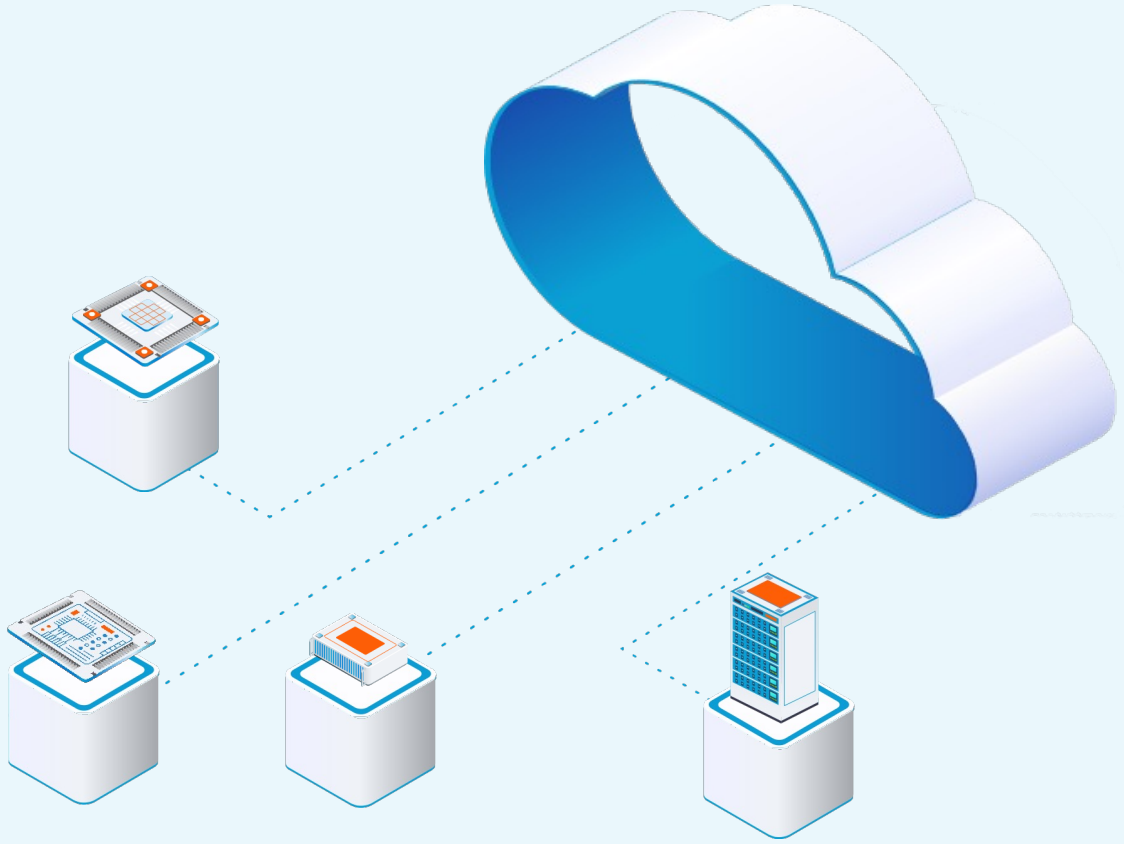

proteanTecs

A New Way of Measuring Heterogeneous SoC Performance

Hanan Moller– Technical Director

Tessent Embedded Analytics - Siemens

Hanan.Moller@siemens.com

Linley Fall Processor Conference

November 1 - 2, 2022

SIEMENS

Agenda

Introduce a performance measurement methodology that combines debug and monitoring IP

Hardware, software and use cases all contributing to complexity

Workloads defined by diverse businesses…

Cloud/ datacenter	Enterprise	Service provider
SMB/ SOHO	Medical & industrial	Automotive

are executed in diverse usage environments…

...using new generations of semiconductor systems

CPU subsystems	Accelerators	Wired interconnect
Interconnect (NoC, inter-chip…)		
Memory	Security	Wireless

Need to measure pre- and post-silicon performance and provide feedback from the field into both dev ops and design for the next generation

SIEMENS

Complexity, increasing software workloads and continuous improvement require a new approach

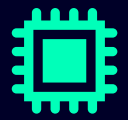

Validation now an essential step in design implementation

- Pre- and post-silicon
- Requires “whole system” view including the real software workload

Existing tools and techniques are inadequate

- Not built for today’s levels of complexity
- Not built for automated testing processes

Feedback of behavioral data is essential

- Move from waterfall to continuous improvement
- Ideally need to retain visibility even after deployment

Traditional emulation, prototype and simulation flow

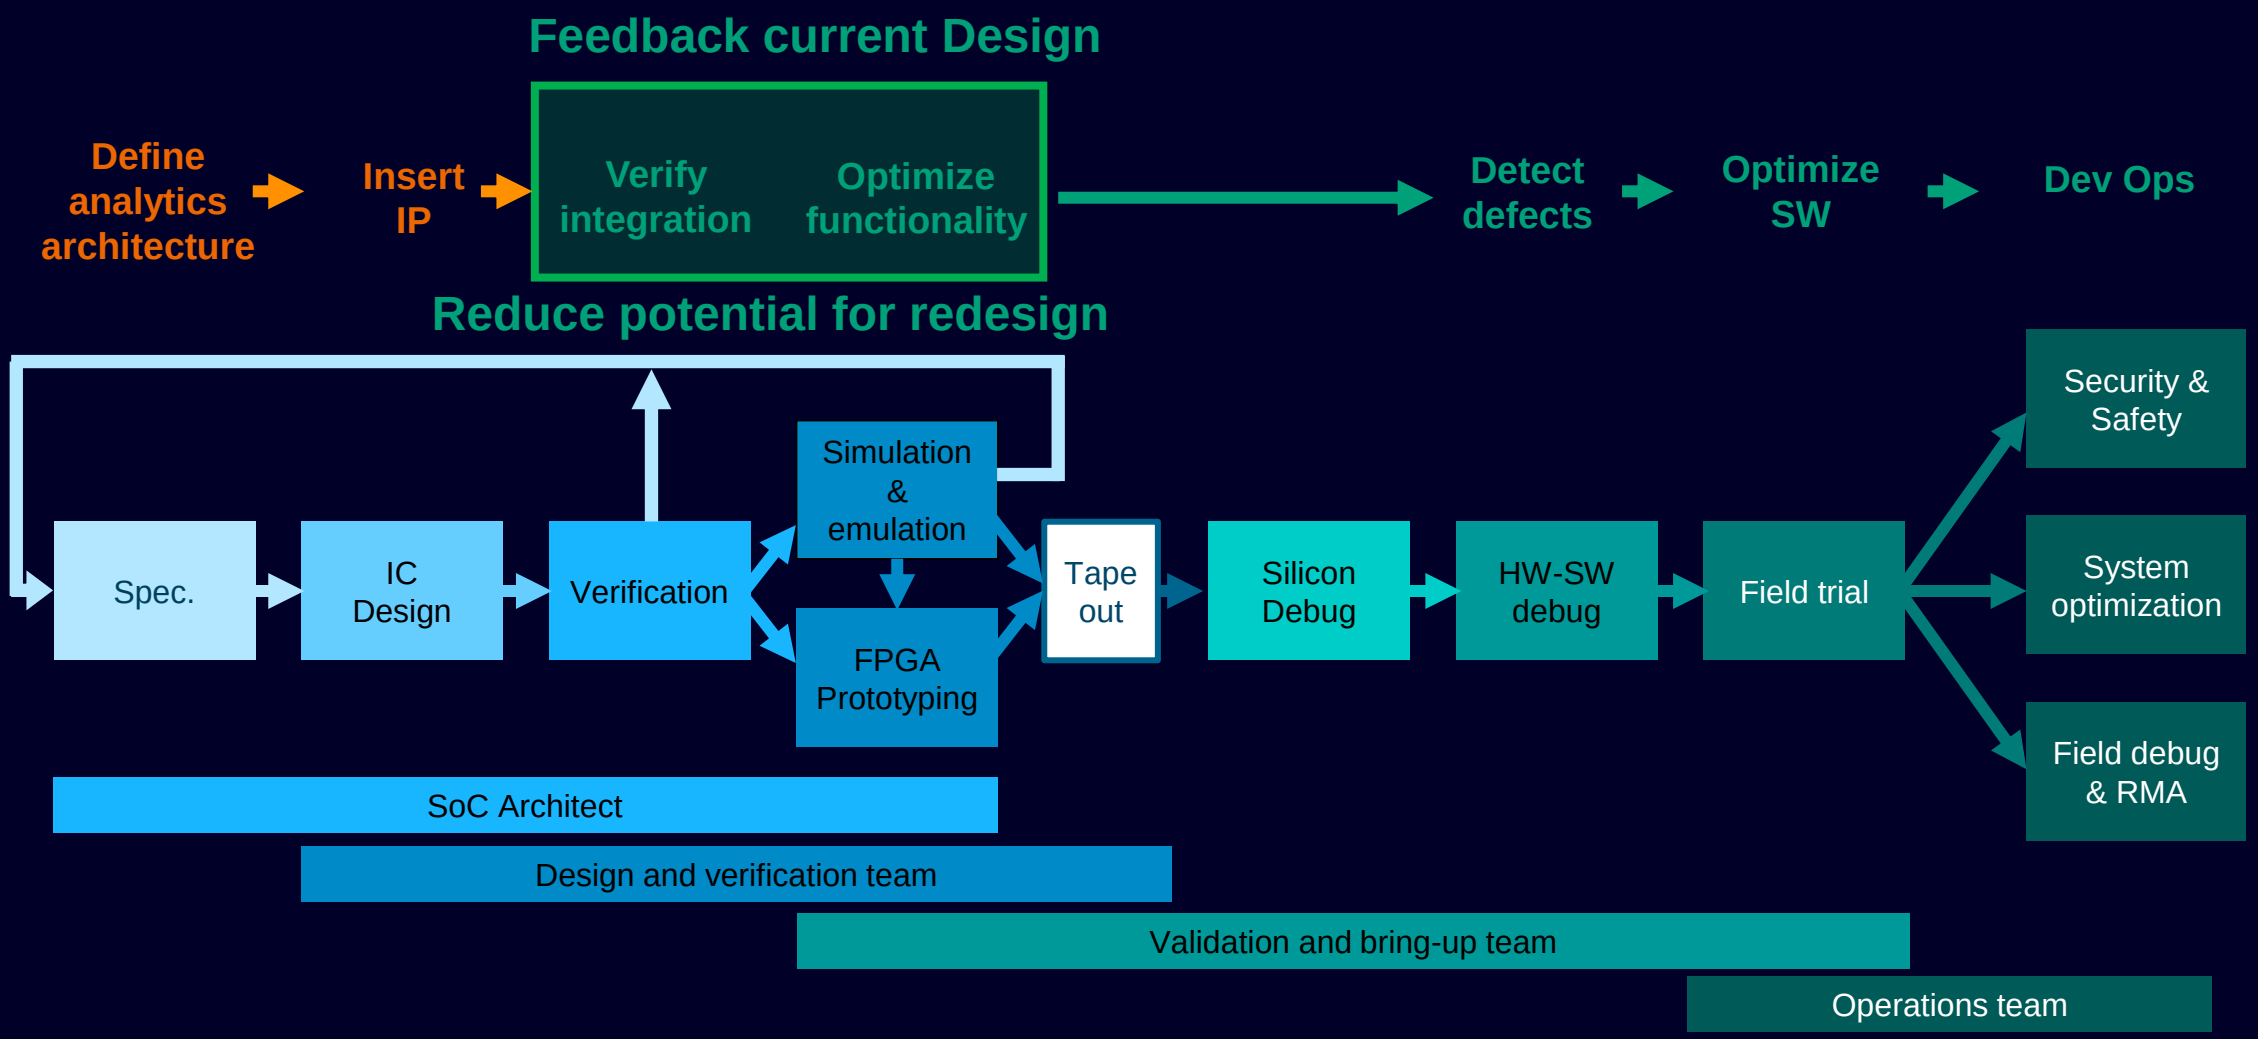

SIEMENS

Adding post-silicon visibility in your development flow

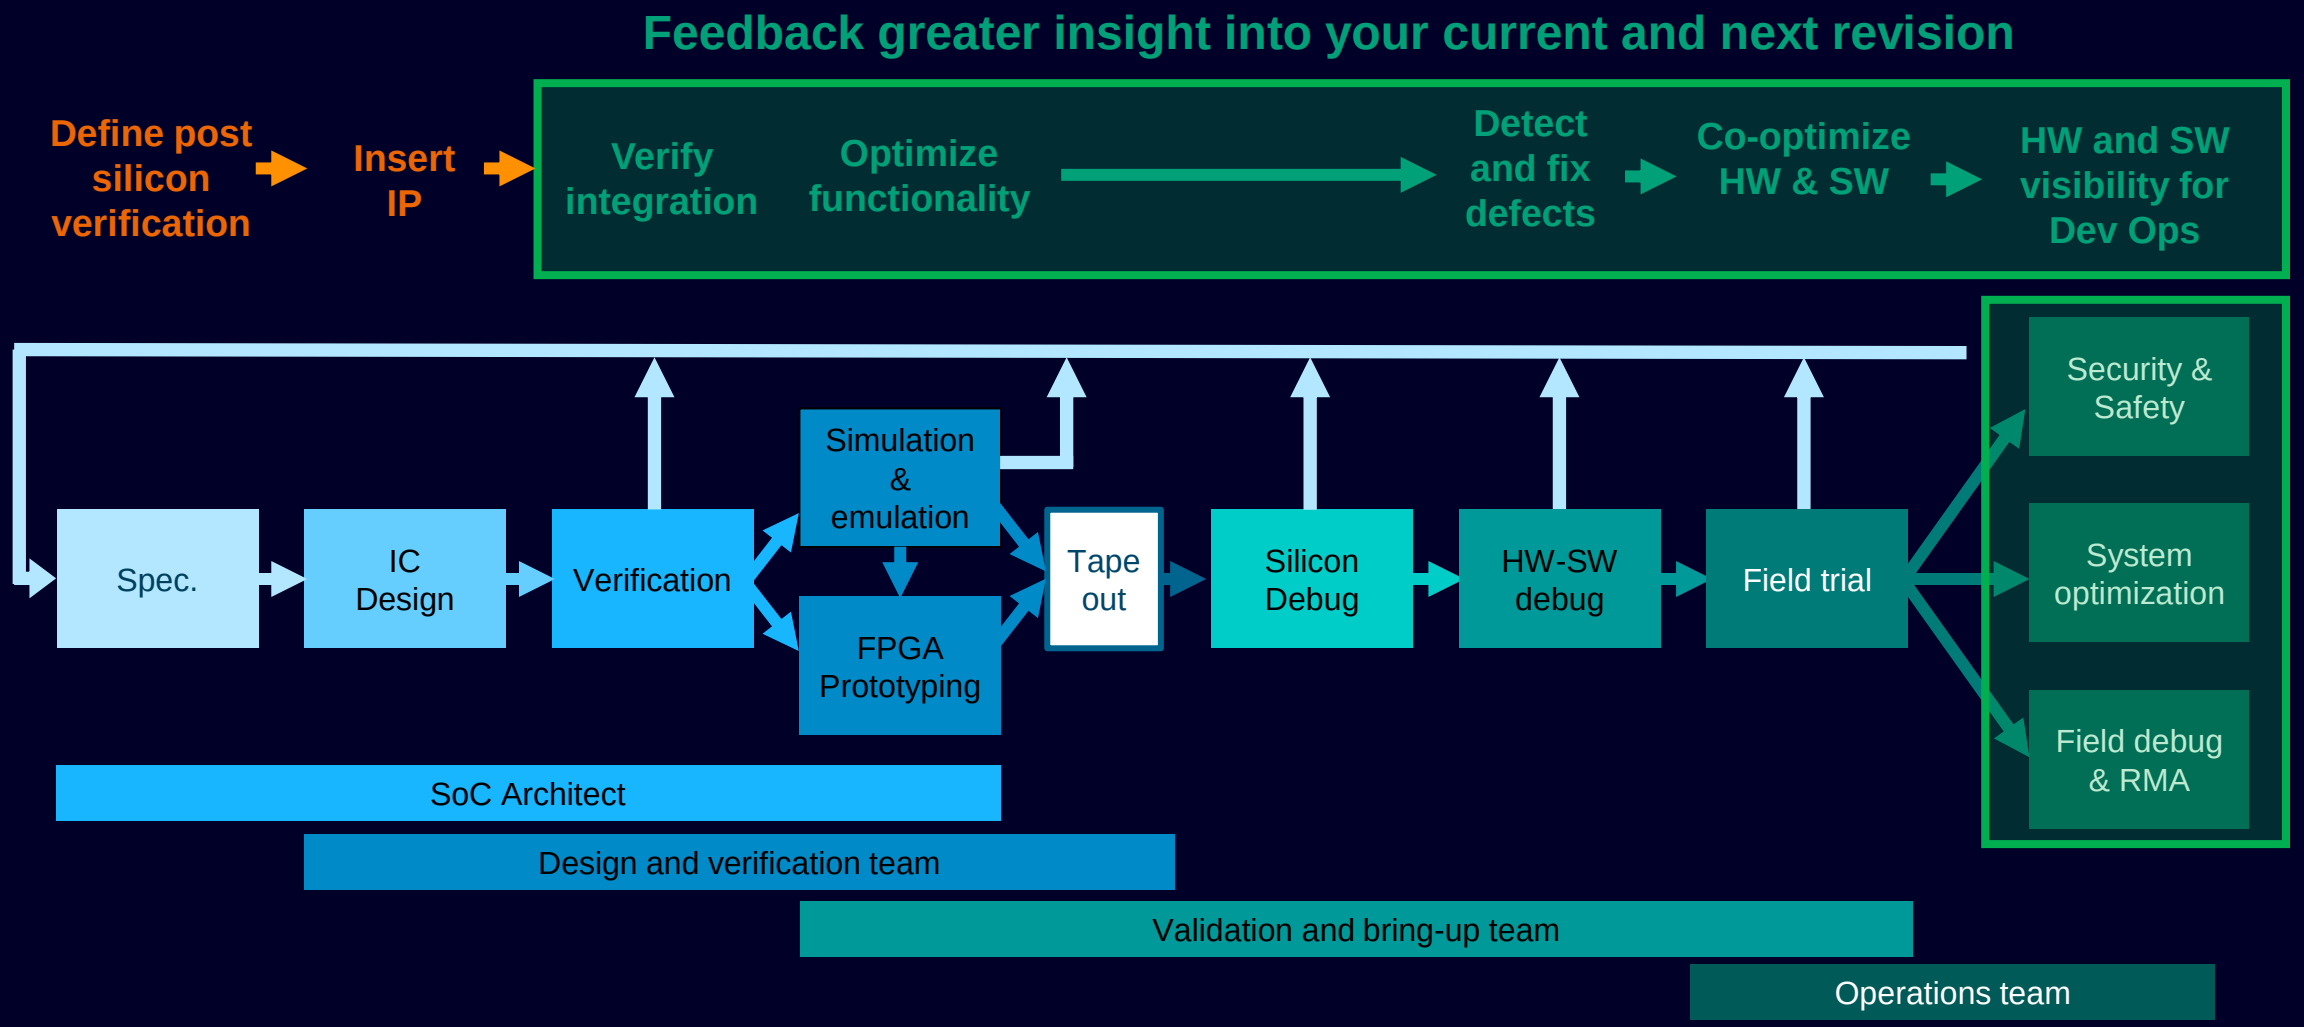

SIEMENS

Consider validation and performance requirements beyond the device

Chip developer

Integrates Embedded Analytics into ICs

EA-enabled devices

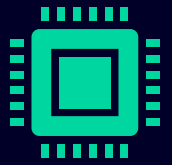

Benefits

- Faster product sign-off
- Fewer bugs
- Effective use of HW resources

Embedded system manufacturer

Builds and monitors products using EA-enabled devices

EA-enabled product fleet

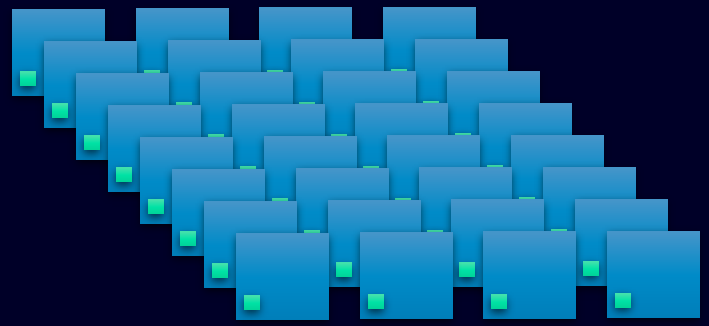

Benefits

- Optimized apps under real workloads
- Monitor safety / secure transactions
- Better remote debug and long tail bug analysis

System operator

Operates products with devices from many manufacturers

EA-enabled facilities

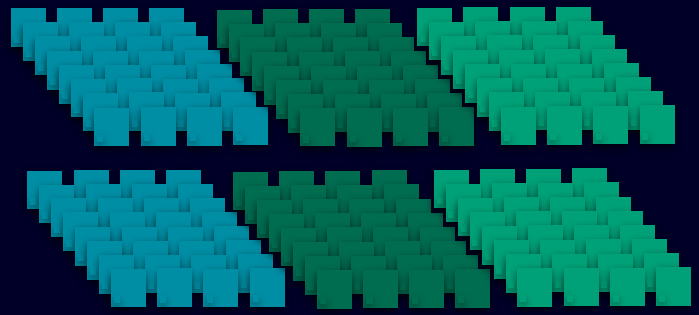

Benefits

- Insights into CAPEX, OPEX, SLAs and security
- RMA

SIEMENS

Common requirements for an automated post-silicon validation solution

Software / firmware teams need hardware designs and software infrastructure that support the move to automated, scalable validation and root cause analysis

- **Concurrent software debug and silicon validation capturing HW/SW interactions**
- **The right data must be observable and available for on-chip and off-chip analysis**
- **The ability to inject, detect or correct faults, errors and other stimulus**
- **Must be able to locally deal with unexpected or non-deterministic behaviors**
- **Remote access to targeted insights and actions, not petabytes of raw data**
- **Instrumentation costs must be kept to a minimum (PPA, BW, servicing etc)**

SIEMENS

Example case study with post-silicon validation

SIEMENS

Example heterogeneous SoC architecture

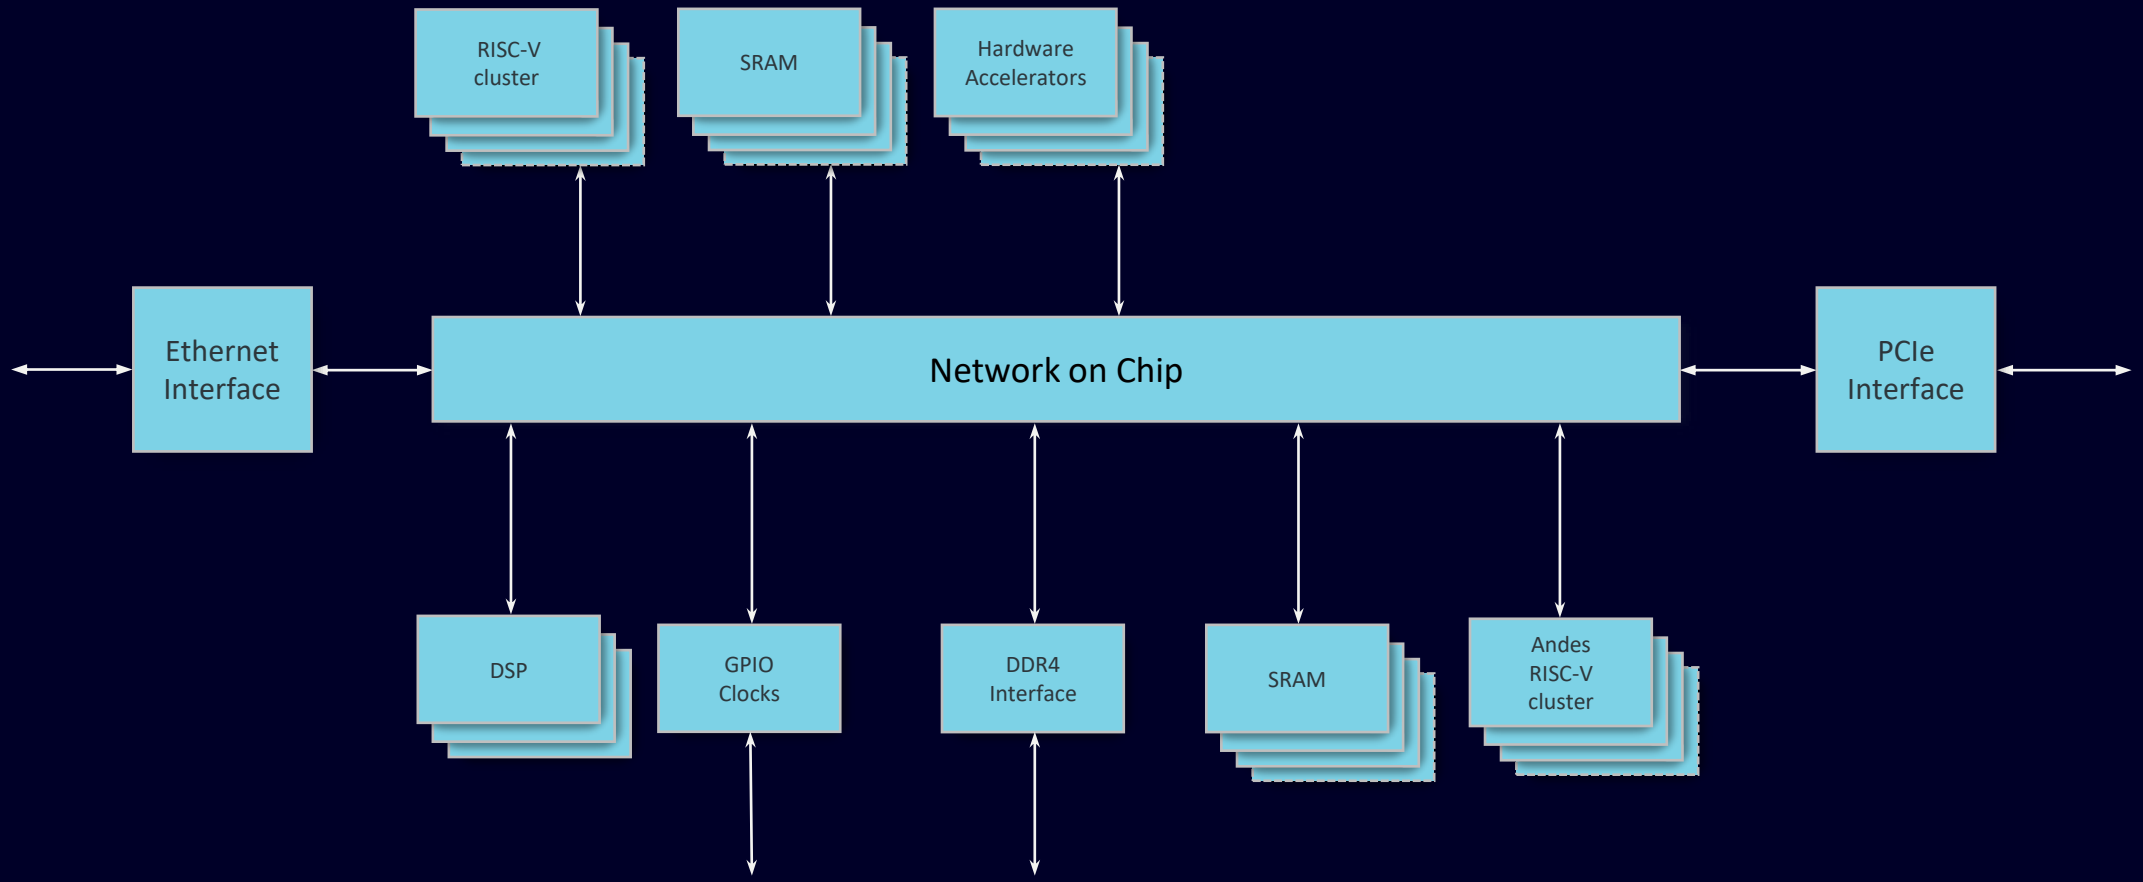

SoC instrumented with validation IP

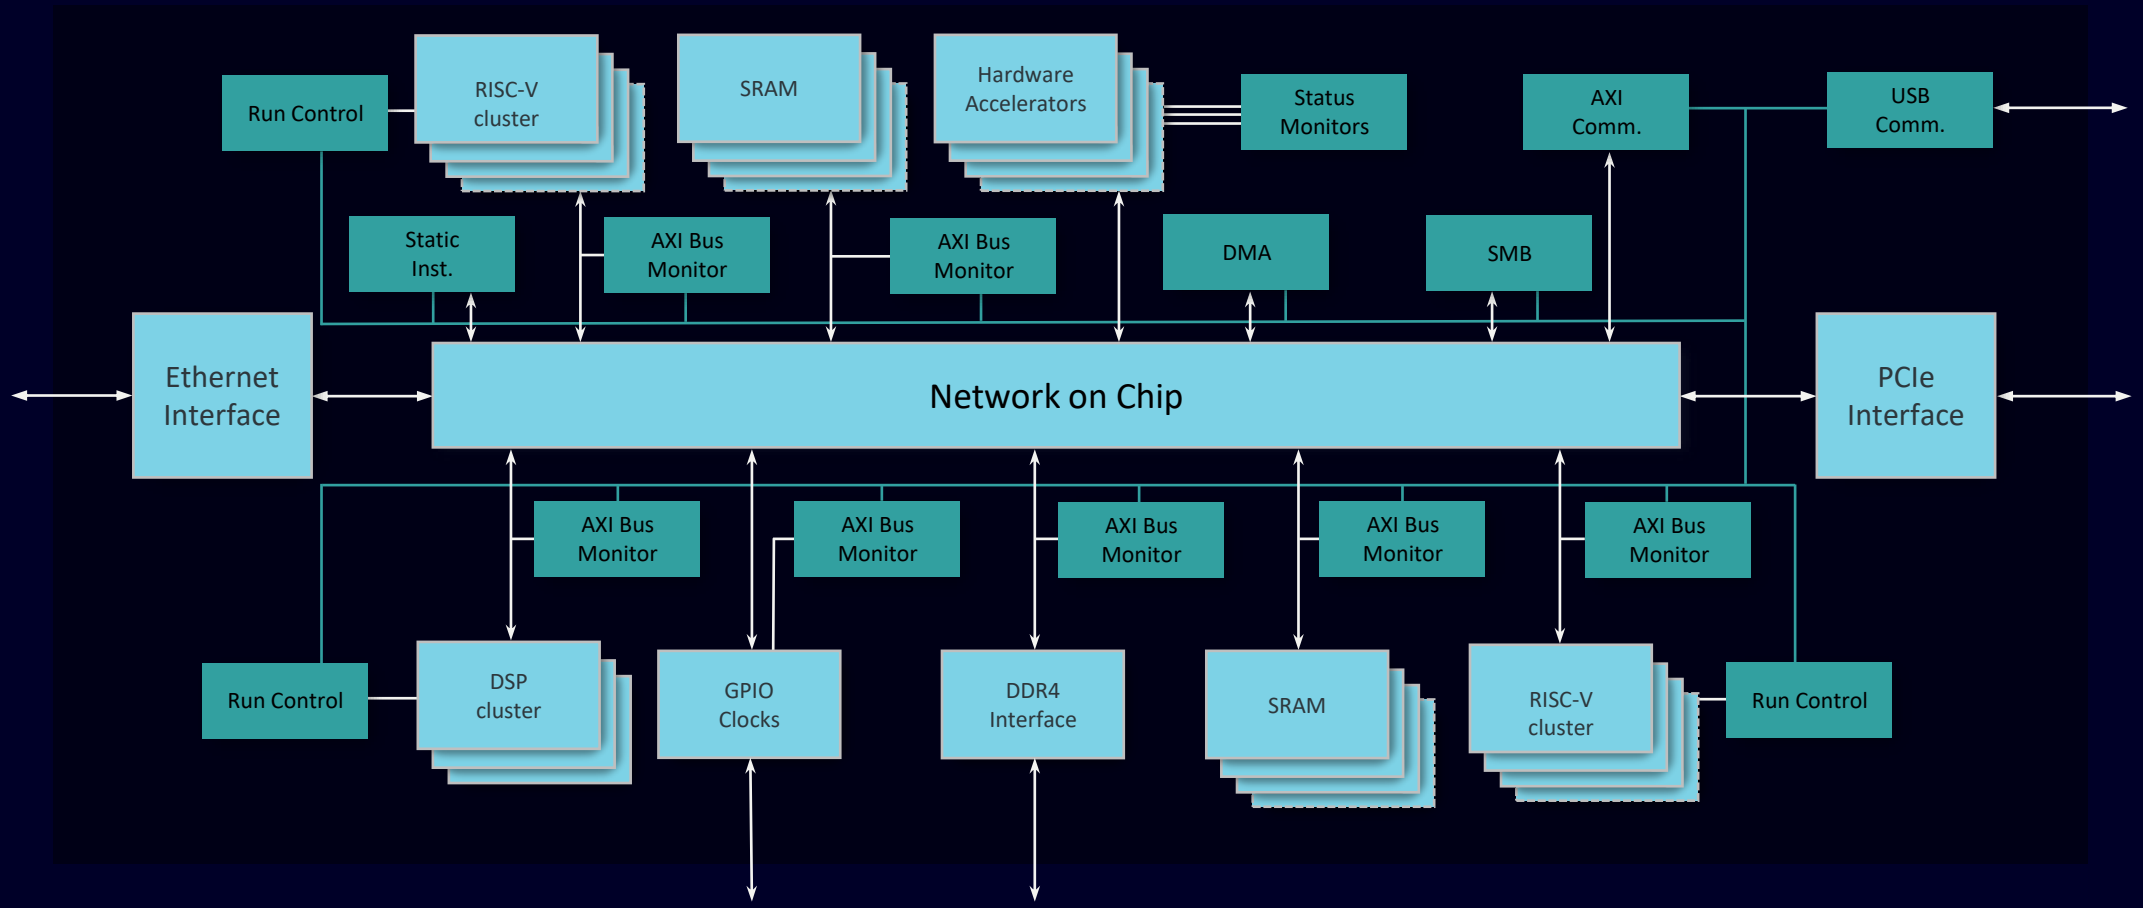

SIEMENS

NOC deadlock detection, out of range - forensic trace

Bus Monitor –
Continuously trace all traffic

Conditional trace triggered if transaction duration exceeds threshold (eg 5k cycles, value or address range)

Signal invalid operation and save trace for analysis

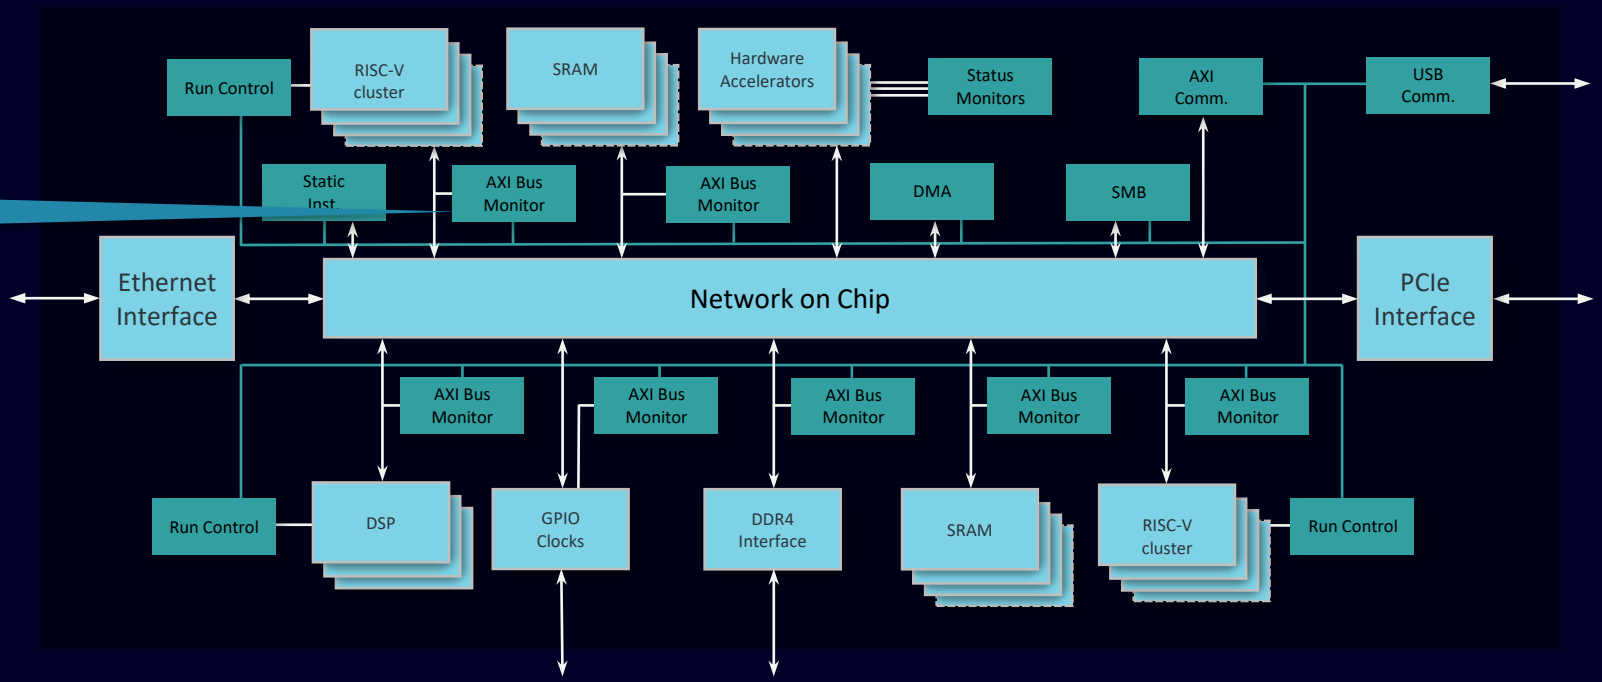

SIEMENS

Application thread timing and CPU performance loss

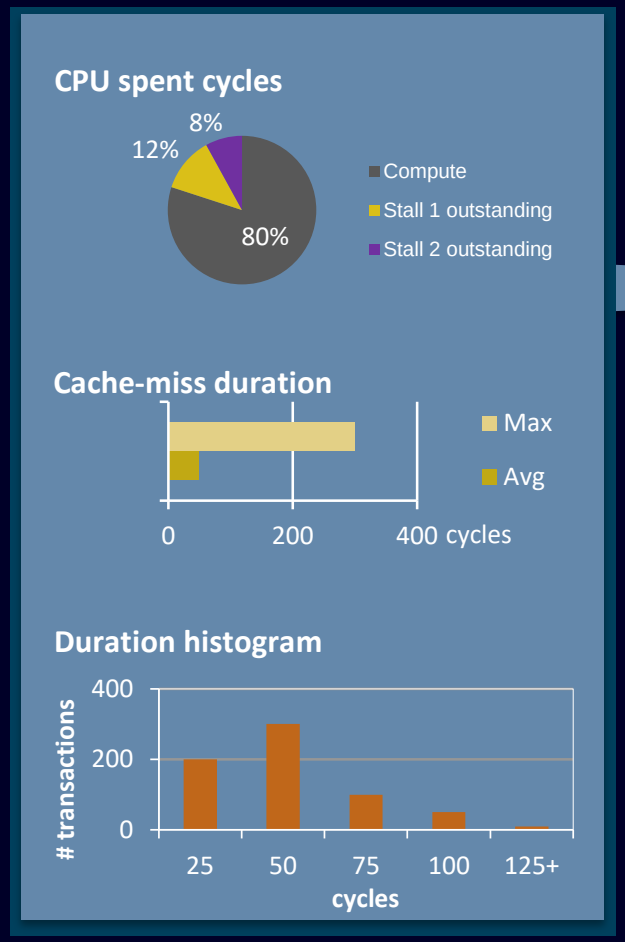

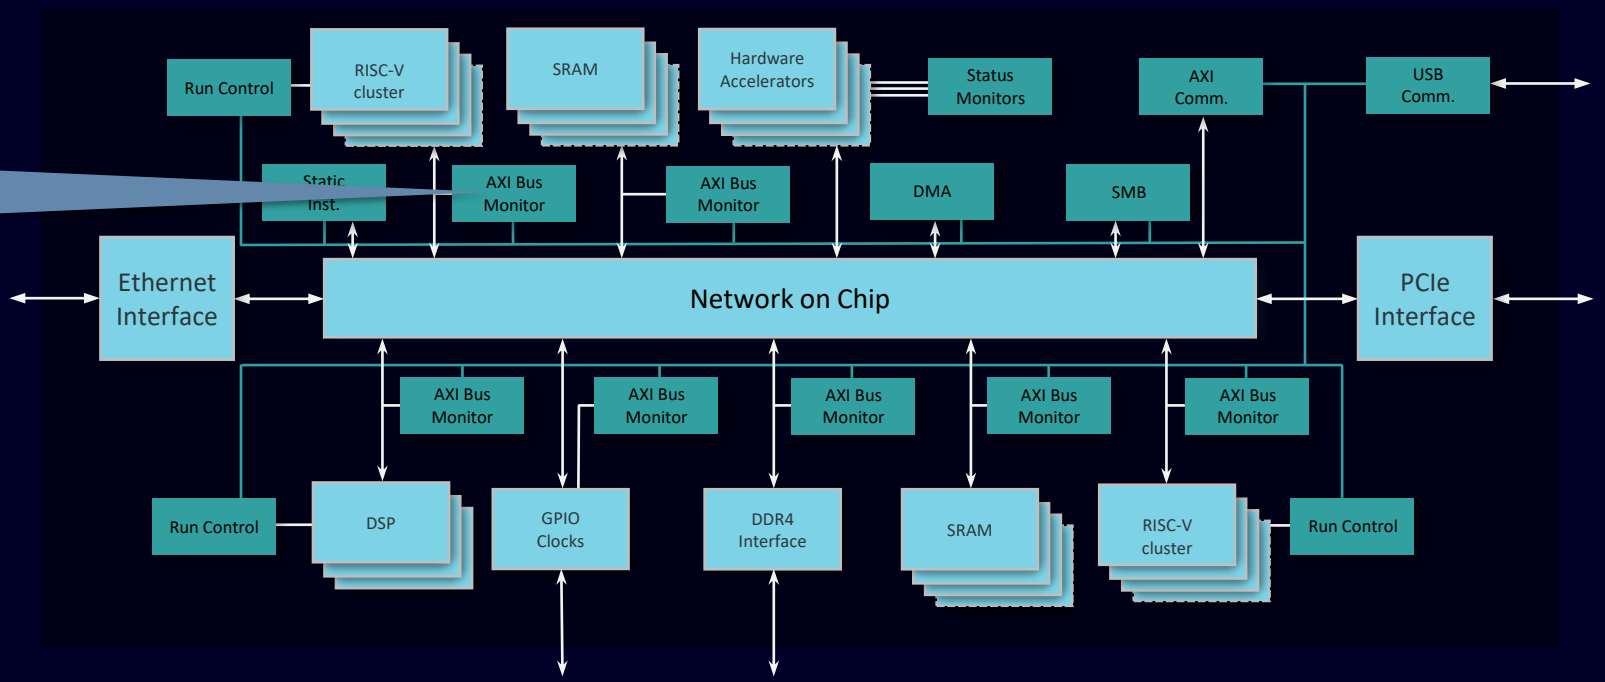

Page 13

Memory bandwidth and access performance

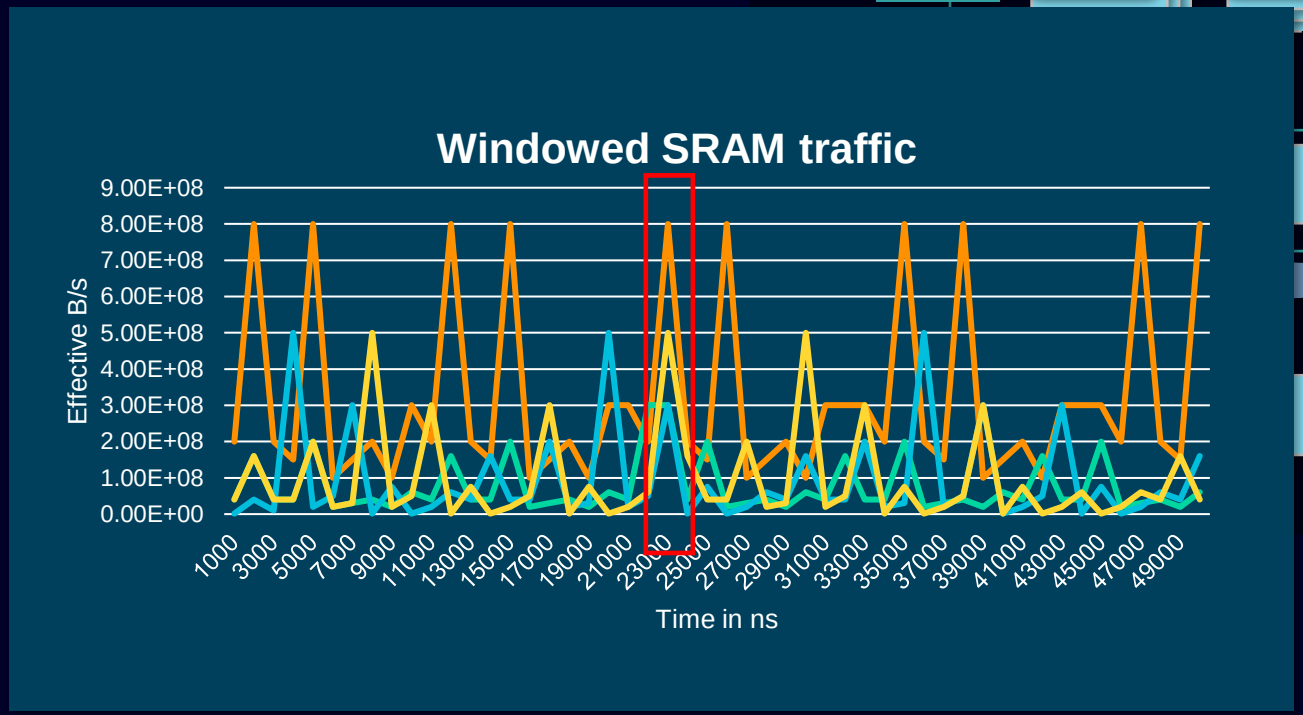

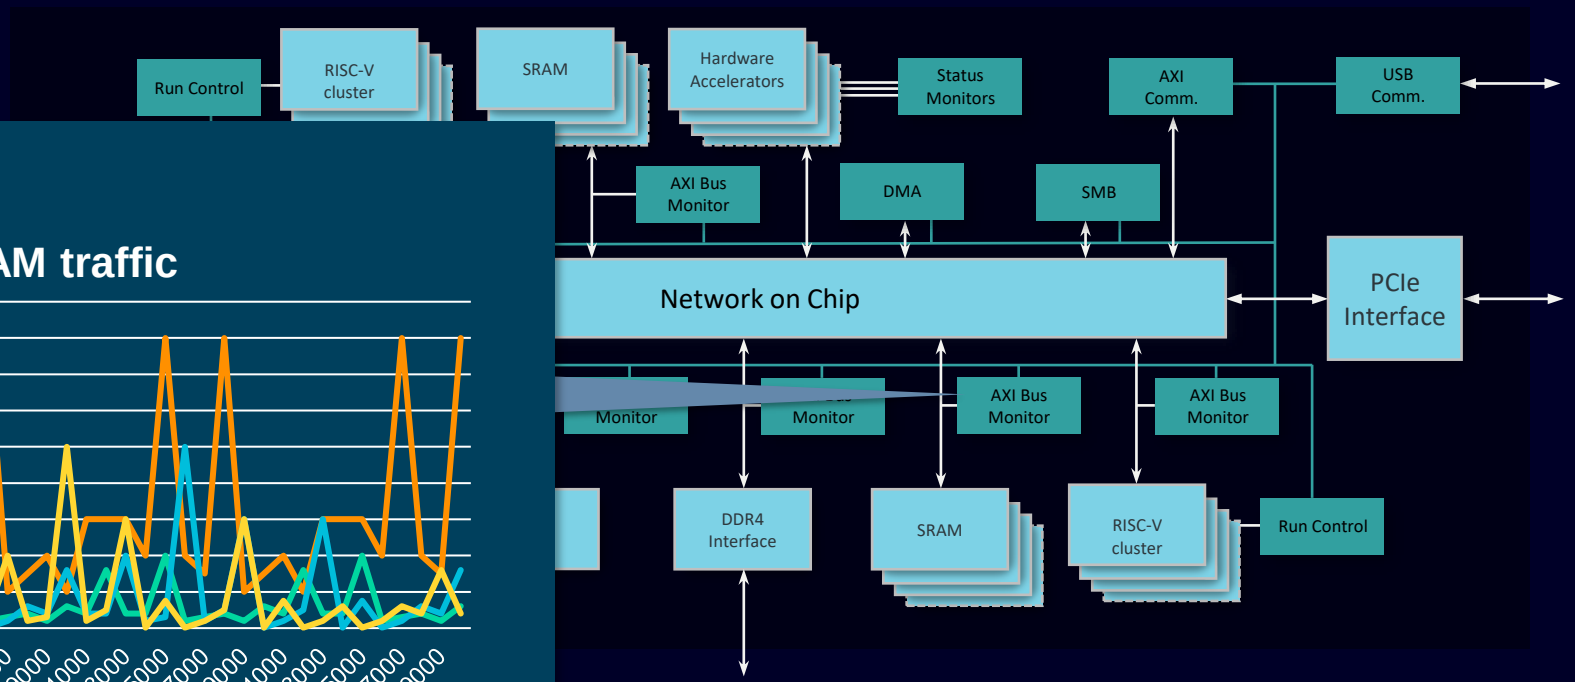

Page 14

SIEMENS

Implementing runtime analytics and reporting on-chip

Use a main system CPU or dedicated processor to configure the monitors and analyze data

Analytics CPU configures monitors and collects data via the AXI Communicator

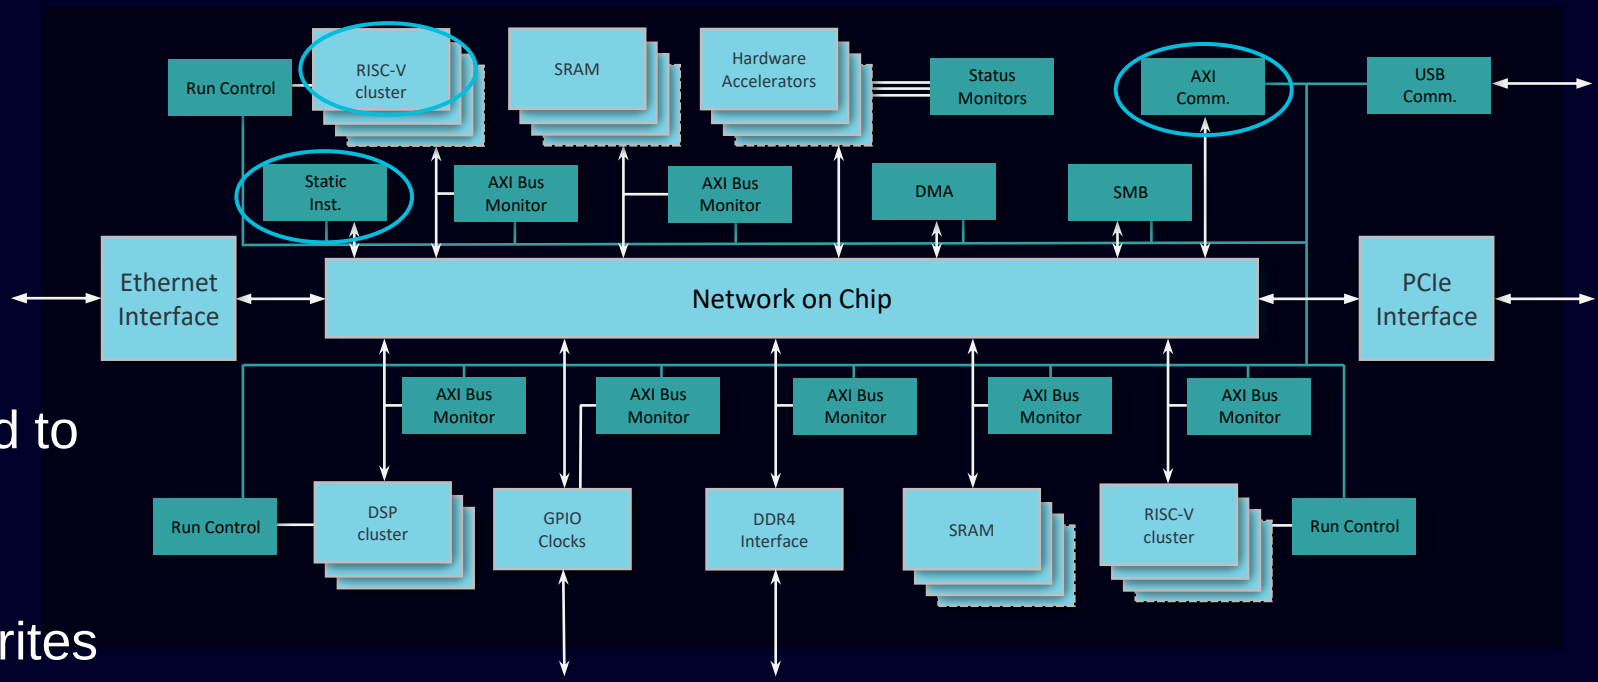

Static Instrumentation module used to determine control flow and timing

Application SW and HW accelerators perform posted writes to here

SIEMENS

A complete technology stack is required for silicon and system lifecycle management – turning 1's and 0's into insight and performance metrics

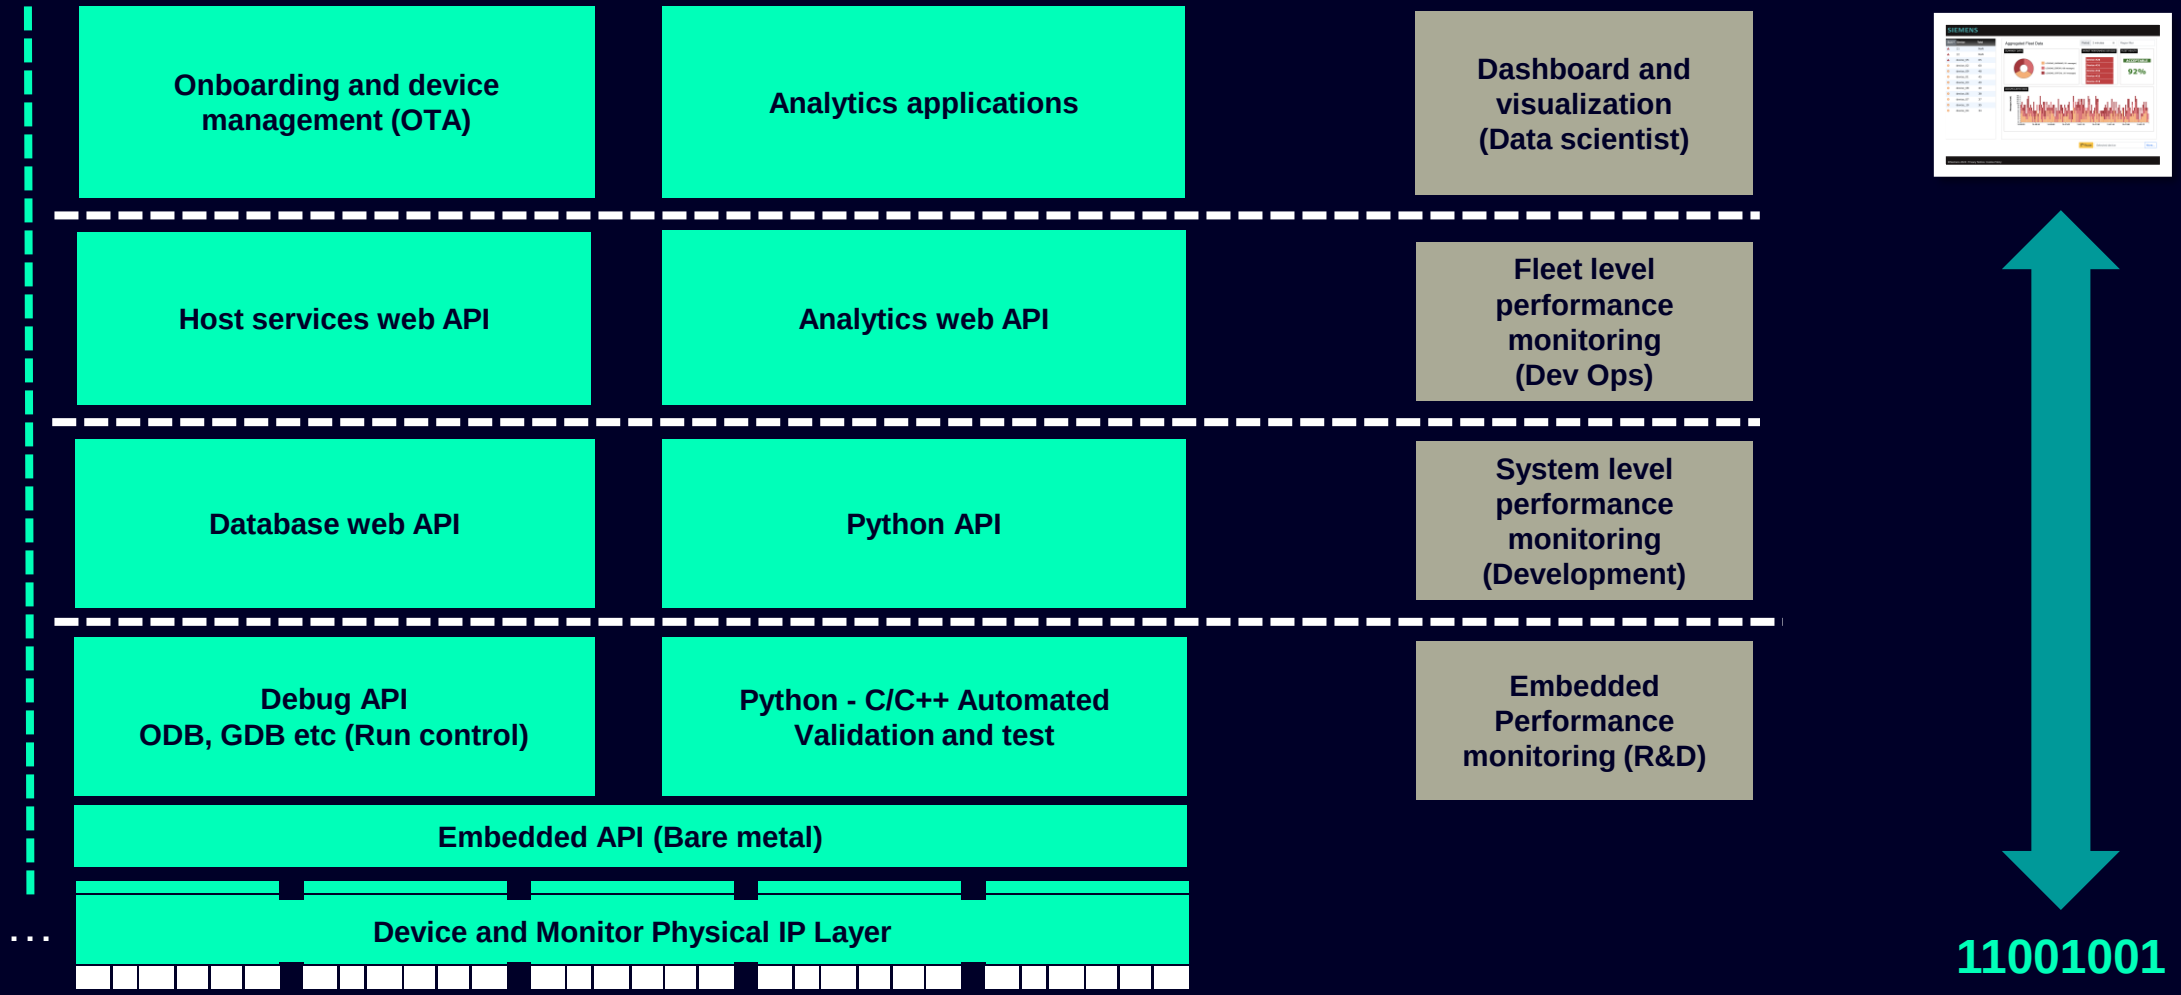

SIEMENS

Conclusions and call(s) to action

- New pressures and working methods mean that traditional debug and validation techniques sometimes aren't enough
  - Increasing complexity | changing workloads | continuous improvement
- Silicon architects and hardware designers: think about system and software validation up-front
  - Technology exists to dramatically improve these processes with minimal PPA penalty
- Software teams: software-based tools aren't everything… the HW team can probably help
  - Support from the hardware team will make your life much easier – and efficient
- Tessent Embedded Analytics provides unprecedented system-level visibility
  - Integrates into automated workflows
  - Automates configuration and data collection
  - Helps identify issues and root-cause them fast – including hard-to-replicate bugs

Contact

Published by Siemens DI SW

Hannan Moller
Technical Director, Solution architecture, Embedded Analytics
Tessent Silicon Lifecycle Solutions

E-mail Hanan.Moller@siemens.com

SIEMENS

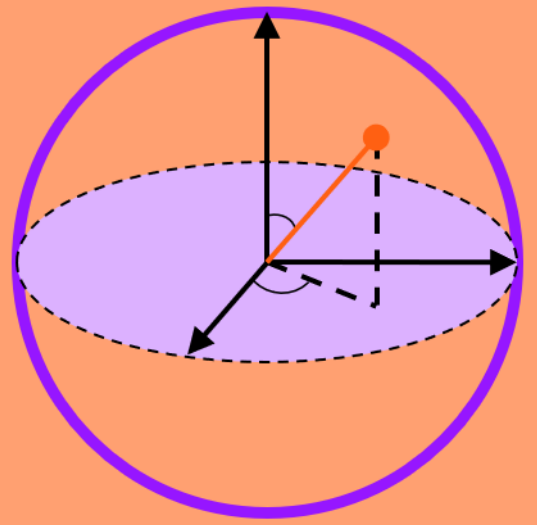

Advanced Semiconductor Manufacturing Technologies for Quantum Computing

Anthony Yu
Vice President, Product Management
GlobalFoundries

anthony.yu@gf.com
Linley Fall Processor Conference, November 1-2, 2022

Foundry support for creating quantum computers

Moving from *bits* to *qubits*

Conventional Computer
0 1
bits

Quantum Computer
|0> |1>
qubits
Quantum state of a two-level system

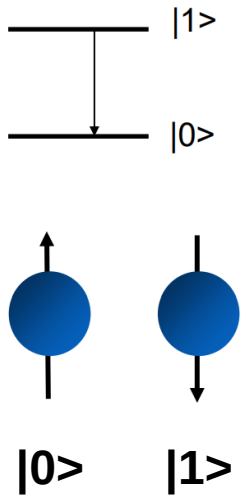

Requires a long term vision to create critical building blocks, whether the approach is matter-based (e.g. quantum dots) or photonics-based

GF platforms enable quantum systems

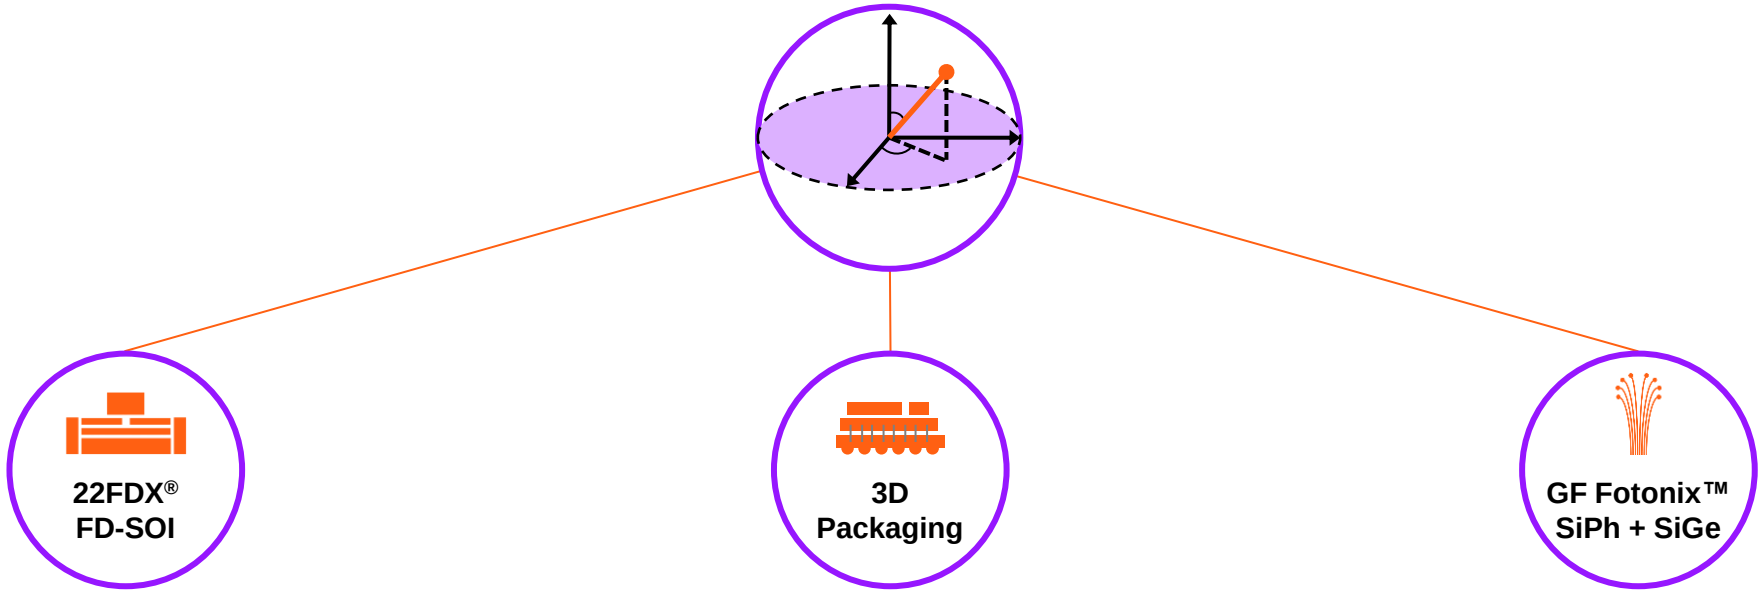

- **Fully Depleted SOI**
- 22nm monolithic integration for IoT / ATV, elite wireless perf
- ULL, sub THz, back gate IP
- Cyrogenic – low temp design
- Spin Qubit CNOT – QD confinement and control in UT SOI

- **3D system integration**
- 12nm/22nm logic and analog with 3D interconnect
- 1.44um 3D pitch W2W/D2W
- Si, FDX, SiPh
- TSV and µBump
- Co-packaged thermal and stress solutions

- **Monolithic SiPh**
- 45nm electronics and photonics in a single chip
- 300GHz ft/fmax, photonic waveguides VIS to Cband
- Integration: laser, fiber, MEMs
- HP SiGe companion tech for LNA, TIA

22FDX® platform offers unique capabilities for QD storage & control

GF's 22nm Fully-Depleted Silicon-on-Insulator technology – Optimal combination of digital and mmWave

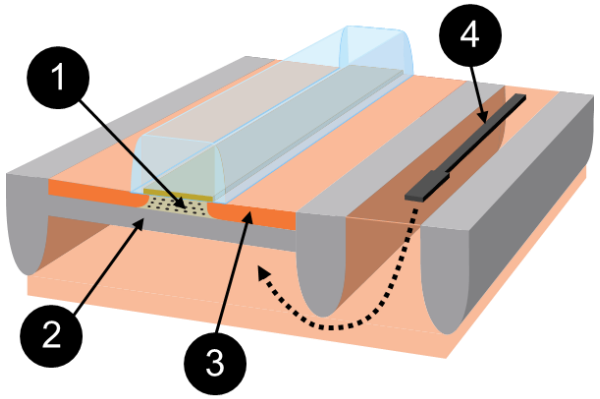

22nm FD-SOI transistor architecture uniquely benefits cryogenic operation

1. Ultra-thin 6nm silicon film allows for small transistor feature size in planar CMOS
   → Provides both transistor speed and spatial confinement

2. Thin 20nm buried oxide layer allows the transistor channel to be undoped
   → Avoids degradation from carrier freeze-out

3. Transistor body is electrically isolated from the channel, but electro-statically influences the channel
   → Body bias as an extra design knob to achieve operating point at cryo temperatures

4. Transistor drain on buried oxide minimizes parasitic capacitance

22FDX® platform for cryo temperature quantum computing

- Key characteristics, such as fT/fMax, improve at cryogenic temperature in 22FDX technology
- Advanced-node minimum feature size helps create higher confinement for quantum dots
- Demonstrated possibility of monolithic quantum processor
- TIA gain measurement of 80dBΩ at 2K, better than at room temperature
- Back-gate bias eases design effort to maintain sweet-spot operating condition at 2K

- IEEE EDL 2019, U. Toronto
- RFIC 2019, U. Toronto

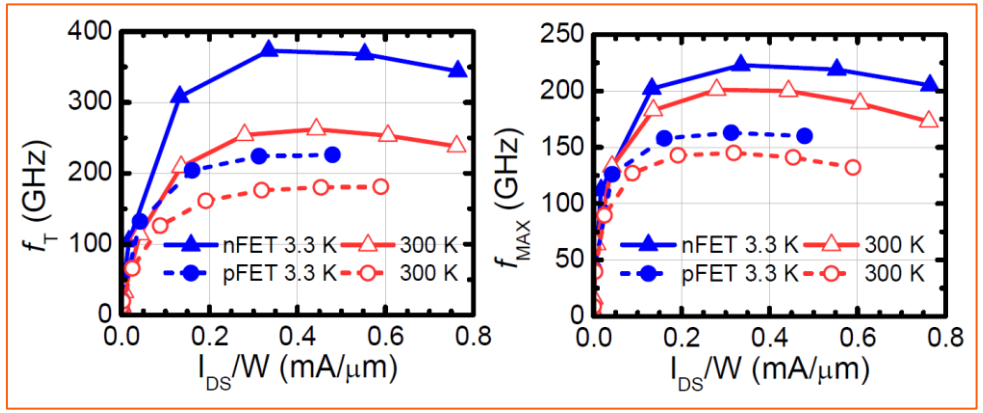

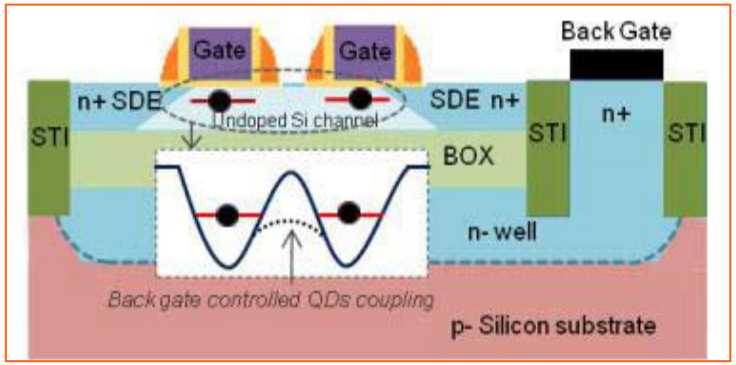

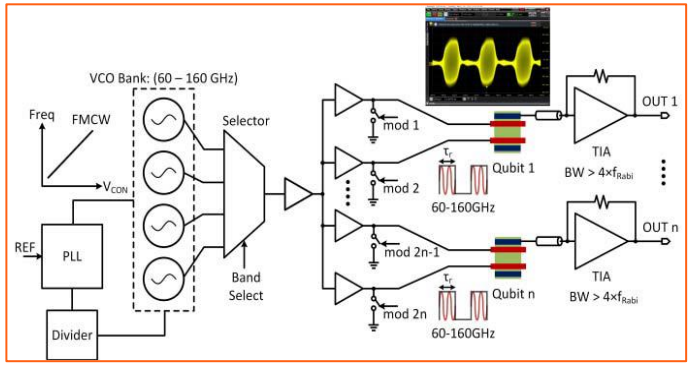

Proof points of 22FDX® platform for quantum computing

Fully integrated quantum computing SOC in 22nm FD-SOI

International Workshop on Cryogenic Electronics for quantum Systems
Fermilab June 19, 2019

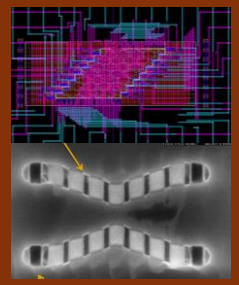

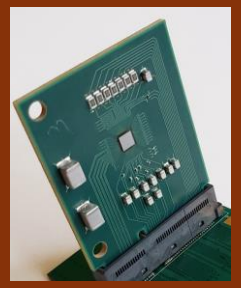

- Full 22FDX SoC to characterize **charge based** qubits
- Fully operational cryo system with control
- Well behaved CMOS characteristics to cryo temps.
- Confirming single electron injection and transport
- Scalable quantum computer under development.

Spin-based Quantum Computing

International Workshop on Cryogenic Electronics for quantum Systems
Fermilab June 19, 2019

- RF designs to control/readout Si **spin based** Qubits.
- Well behaved CMOS characteristics to cryo temps.
- Early quantum dot behavior under evaluation.
- Targeting initial Qubit capability by end 2019.
- Integration and scaling options under development.

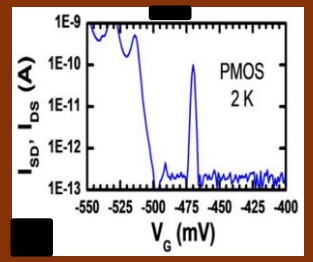

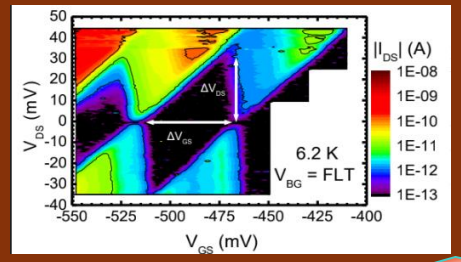

Bias Generation and Calibration of CMOS Charge Qubits at 3.5 Kelvin in 22-nm FDSOI

Imran Bashir*, Dirk Leipold*, Mike Asker*, Ali Esmailiyan‡, Elena Blokhina†‡, David Redmond‡, Panagiotis Giounanlis‡, Dennis Andrade-Miceli†‡, Bogdan Staszewski†‡
*Equal1 Labs, Fremont, CA, USA. † University College Dublin, Dublin 4, Ireland.
‡Equal1, NovaUCD, Belfield Innovation Park, University College Dublin, Dublin 4, Ireland.
E-mail: imran.bashir@ieee.org

ESSCIRC 2021

Cryogenic Compact Low-Power 60GHz Amplifier for Spin Qubit Control in Monolithic Silicon Quantum Processors

Michele Spasaro[1], Shai Bonen[2], Gregory Cooke[2], Thomas Jager[2], Tan D. Nhut[1], Dario Sufrà[1], Sorin P. Voinigescu[2], Domenico Zito[1]
[1]Dept. of Electrical and Computer Engineering, Aarhus University, Denmark
[2]Edward S. Rogers Sr. Dept. of Electrical and Computer Engineering, University of Toronto, Canada
domenico.zito@ece.au.dk

IEEE IMS 2022

A 4.2-9.2GHz Cryogenic Transformer Feedback Low Noise Amplifier with 4.5K Noise Temperature and Noise-Power Matching in 22nm CMOS FDSOI

Boce Lin[1,2], Hamdi Mani[3], Phil Marsh[4], Richard Al Hadi[5], Hua Wang[1,2]
[1]School of ECE, Georgia Tech, USA
[2]ETH Zurich, Switzerland
[3]CryoElec LLC, USA
[4]Carbonics Inc, USA
[5]Alcatera Inc, USA
blin96@gatech.edu, hua.wang@ece.gatech.edu

IEEE RFIC 2022

Cryogenic Compact mm-Wave Broadband SPST Switch in 22nm FDSOI CMOS for Monolithic Quantum Processors

Tan D. Nhut[1], Shai Bonen[2], Gregory Cooke[2], Thomas Jager[2], Michele Spasaro[1], Dario Sufrà[1], Sorin P. Voinigescu[2], Domenico Zito[1]
[1]Dept. of Electrical and Computer Engineering, Aarhus University, Denmark
[2]Edward S. Rogers Sr. Dept. of Electrical and Computer Engineering, University of Toronto, Canada
domenico.zito@ece.au.dk

IEEE IMS 2022

GF platforms enable quantum systems

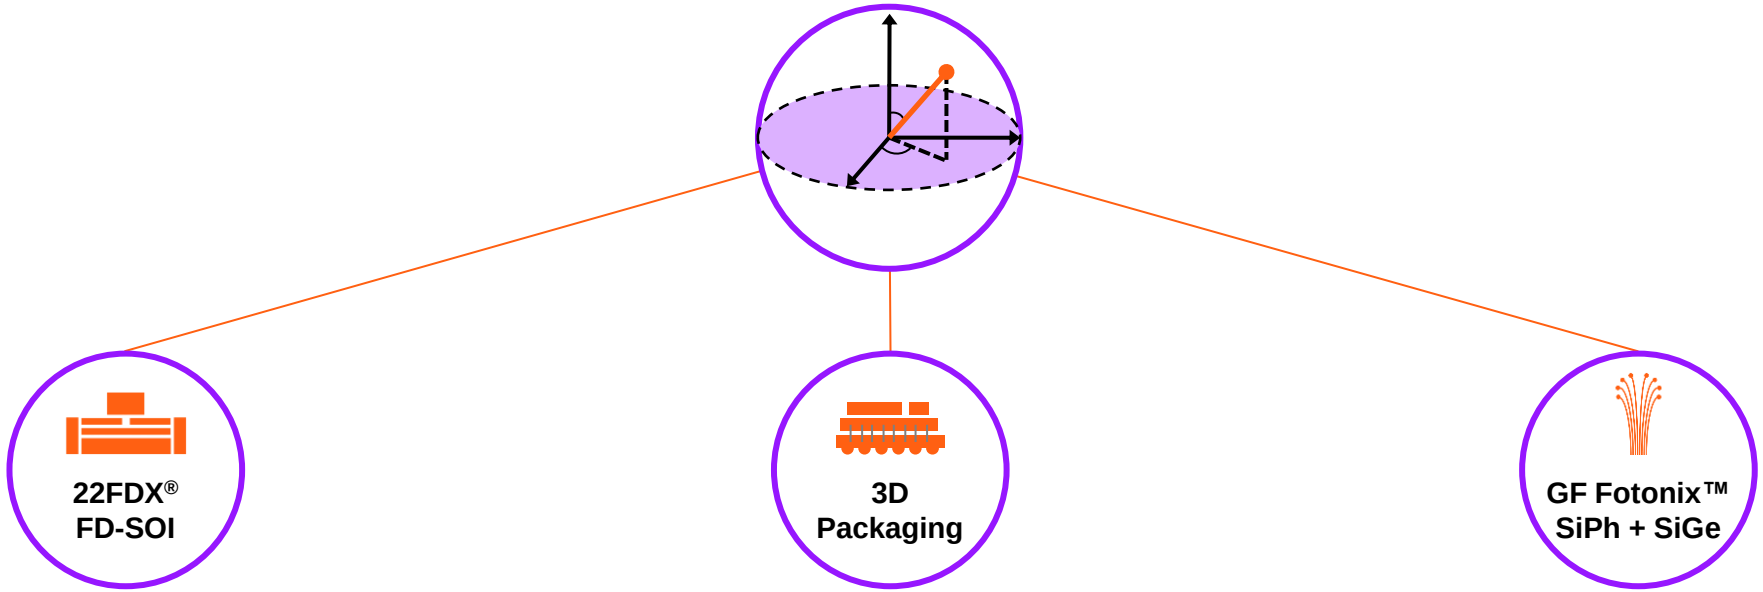

- **Fully Depleted SOI**
- 22nm monolithic integration for IoT / ATV, elite wireless perf
- ULL, sub THz, back gate IP
- Cyrogenic – low temp design
- Spin Qubit CNOT – QD confinement and control in UT SOI

- **3D system integration**
- 12nm/22nm logic and analog with 3D interconnect
- 1.44um 3D pitch W2W/D2W
- Si, FDX, SiPh
- TSV and uBump
- Co-packaged thermal and stress solutions

- **Monolithic SiPh**
- 45nm electronics and photonics in a single chip
- 300GHz ft/fmax, photonic waveguides VIS to Cband
- Integration: laser, fiber, MEMs
- HP SiGe companion tech for LNA, TIA

Many ways to make a qubit

- Qubits can be any two-level system
  - Spin of an electron
  - Polarization of a photon
  - Spin of a quantum dot
- Success of a quantum system depends on:
  - Making
  - Manipulation
  - Entangling
- Need to remain in quantum states of superposition and entanglement (coherence)
  - Coherence times long enough to allow for computation and for error correction

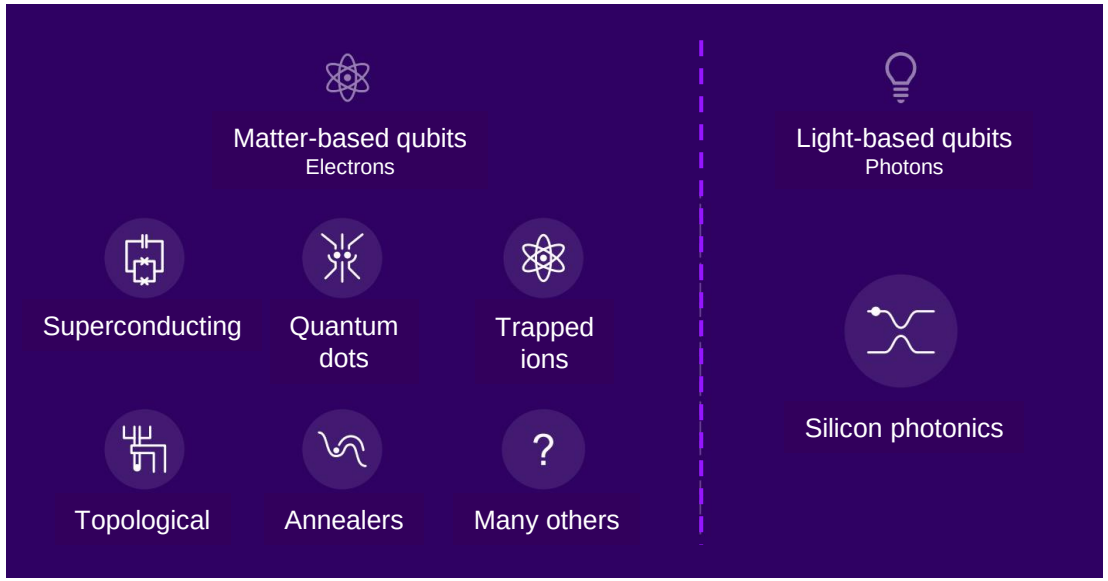

Source: PsiQuantum

Advantages of photons as qubits

- Light (photons)
  - Carries quantum information
  - Works across all temperatures
  - No crosstalk
  - Can be manipulated using optical components built in a CMOS line
- Silicon Photonics is manufacturable

DC optical interconnects underpin growth of photonics foundry market

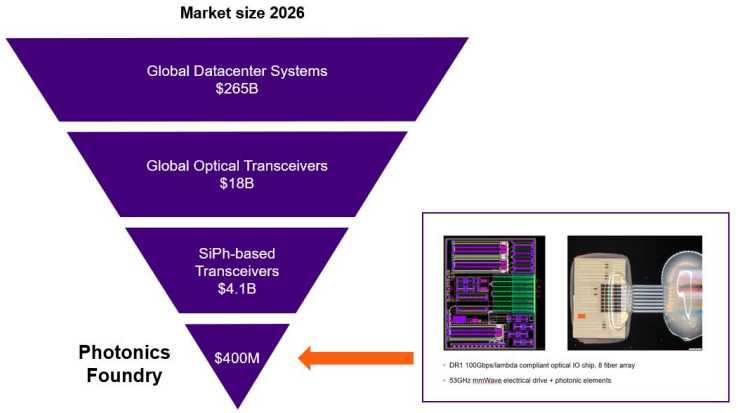

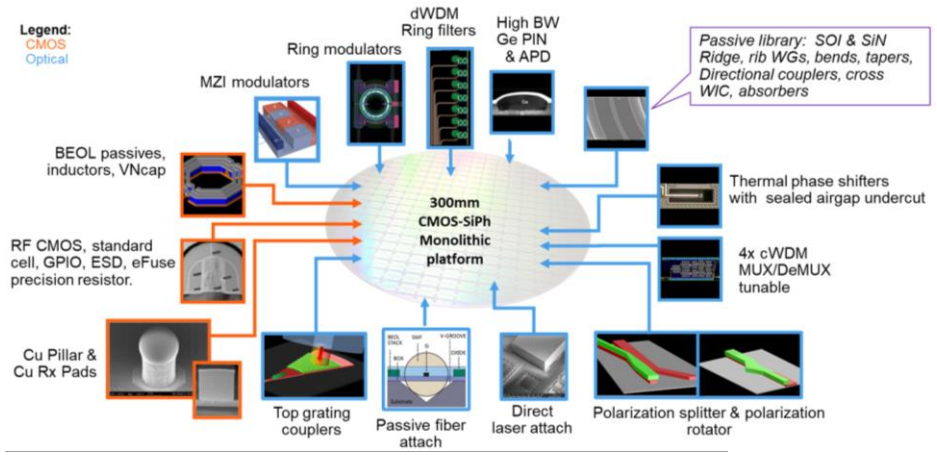

GF Fotonix™ monolithic integration

GF Fotonix™ passive library

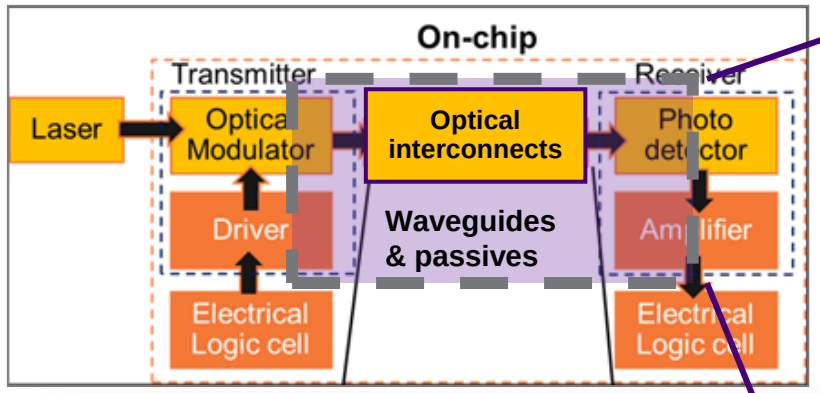

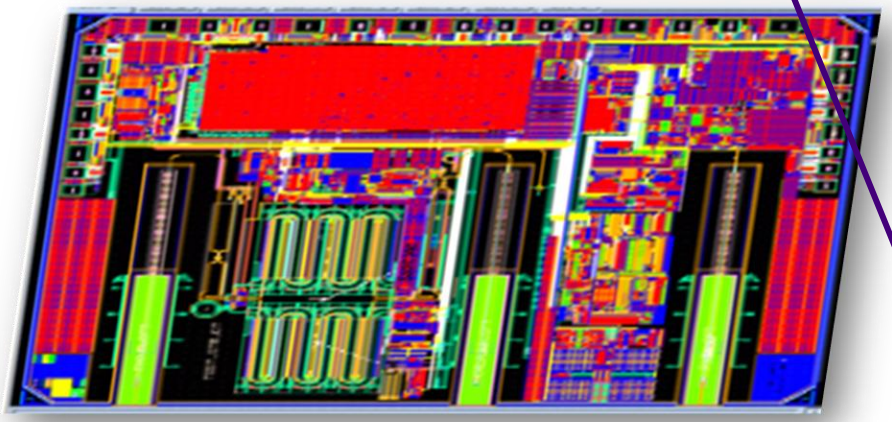

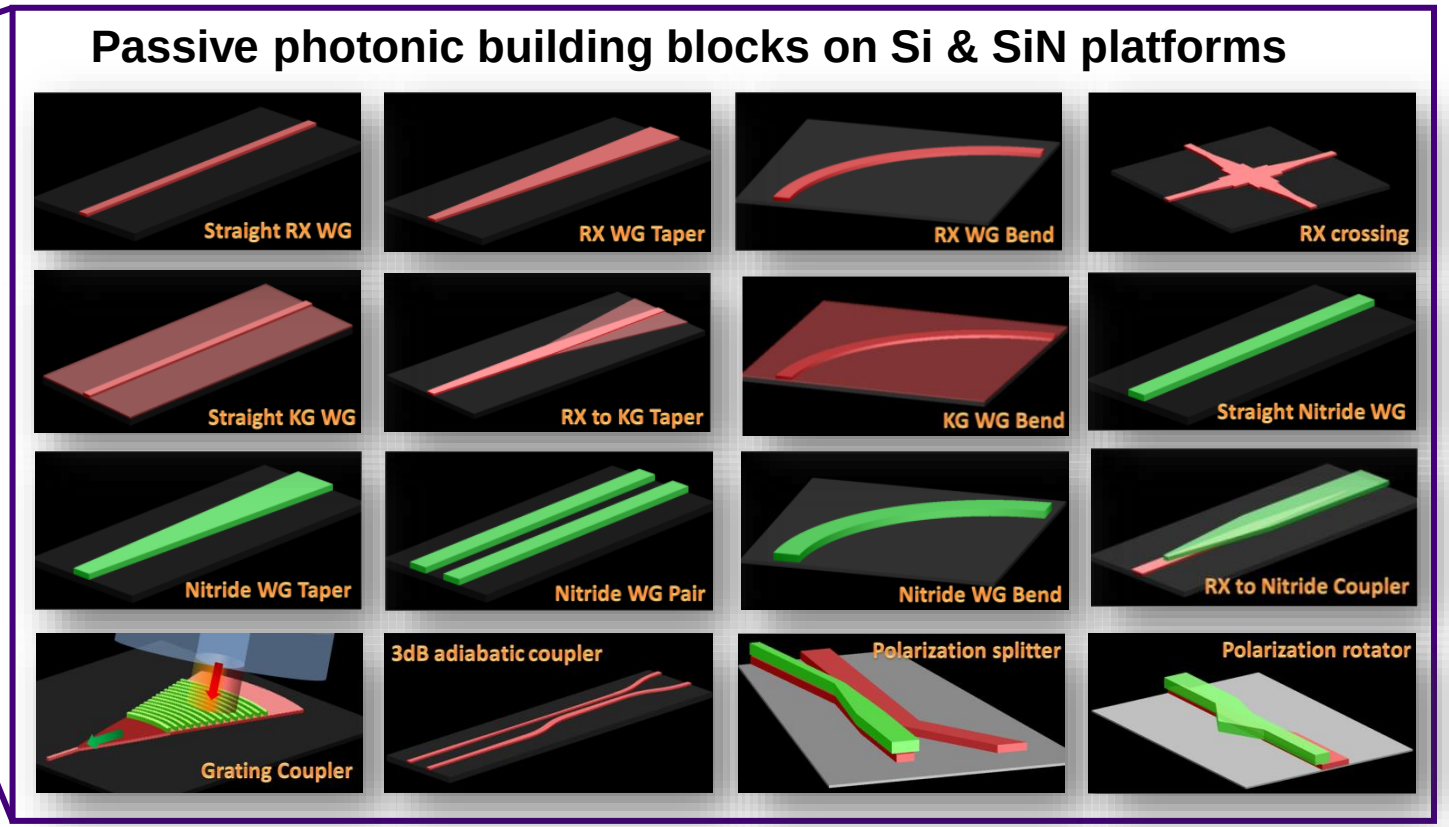

- Routing, coupling, polarization handling, wavelength filtering, etc.
- Si and SiN-based components covering both O and C-bands
- SOI SMF (-1.03dB/cm), SOI Rib (0.73dB/cm)l SiN (0.3dB/cm)

Reference

K. Giewont, et al., IEEE JSTQE, 25(5) 1-11 (2019).
M. Rakowski, et al., OFC, T3H.3 (2020).
Y. Bian, et al., OSA FiO, FW5D.2 (2020).

Building blocks

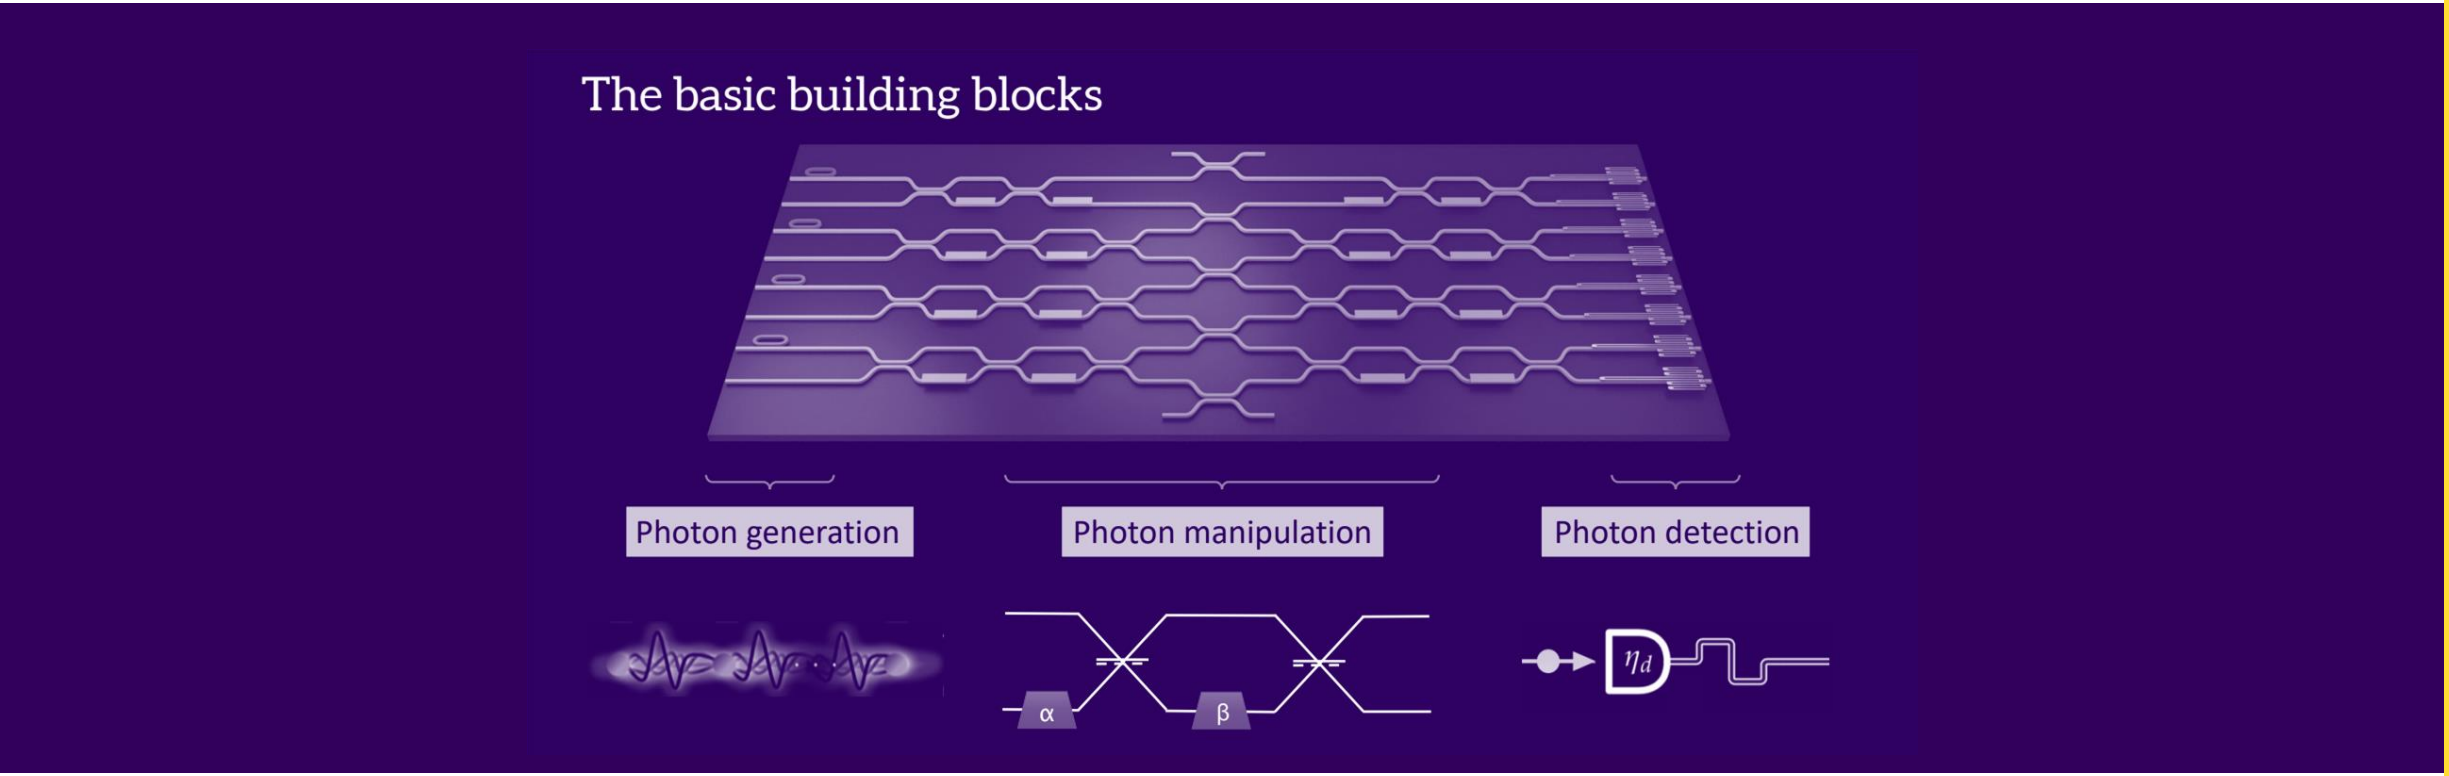

Source: PsiQuantum

Manufacturing for quantum computing

Loss is the enemy: True in all optics, paramount in quantum computing

Datacom/Telecom

- Waveguide loss = <1 dB/cm
  - Supports link budgets for traditional pluggable transceivers or next generation co-packaged optics
  - Silicon waveguides accomplish this in GF Fotonix™
  - Utilizes advanced immersion lithography for reduced Line Edge Roughness (LER)
  - ~2X improvement in loss over non-immersion

	Dry Lithography		Immersion Lithography	
Features	**CD Sigma (nm)**	**LER**	**CD Sigma (nm)**	**LER**
150nm Waveguide	0.99	6.4	0.32	3.5
350nm Waveguide	0.87	6.8	0.40	3.8
400nm Waveguide	0.94	6.1	0.31	3.3

Photonic Quantum Computing

- Waveguide loss = <<1 dB/cm
  - Necessary to support single photon fidelity from generation to detection
  - Utilizes same advanced immersion lithography
  - Incorporates SiN for key low loss waveguide structures
    - Common material, already used in GF FotonixTM
    - Optimizing the deposition process for extremely low loss performance

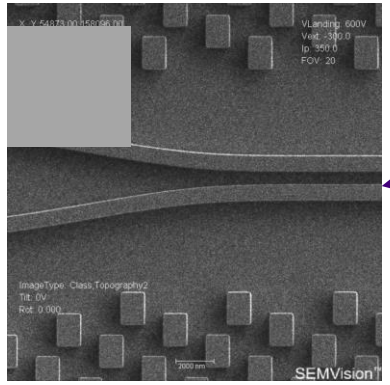

SiN waveguides will support very low loss waveguides (top view)

Anthony Yu, November 1, 2022

Manufacturing for quantum computing

Same base process, new material integration

- Photonic-based quantum computing leverages the advantages of GF Fotonix™ silicon photonics platform
  - Advanced tooling – Industry-leading manufacturing
  - Rigorous process controls – Manufacturability and reliability
- The full life cycle of the photon is supported in silicon
  - Photon creation through detection
  - This makes scaling to a 1,000,000 qubit machine possible
- Quantum computing has specific component needs
  - High speed optical switches – Significantly better loss than Si-based switches
    - Efficient, low loss manipulation of single photons
  - Single photon detectors – Superconducting detection
    - Key basis of qubit creation to drive compute architecture
  - These are being integrated and using our standard process and toolset

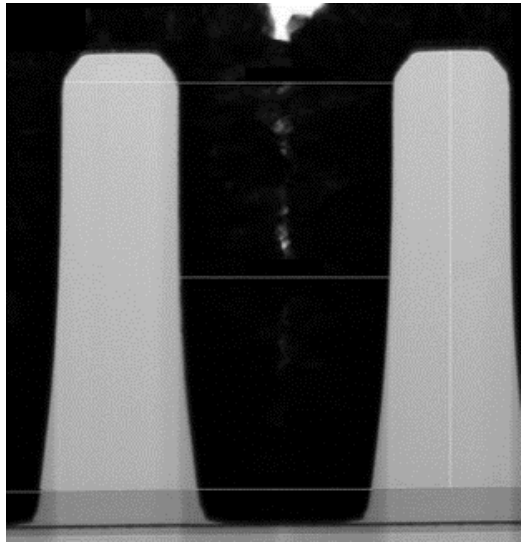

Single photon detector

Manufacturable components built for quantum computing

Manufacturable components

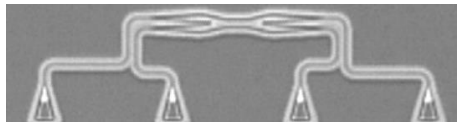

Coupler

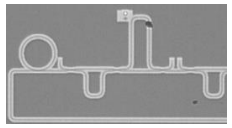

Filter

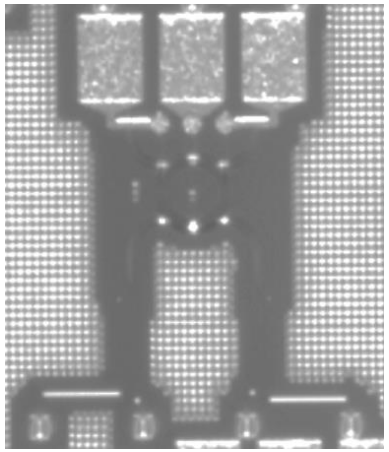

Switch

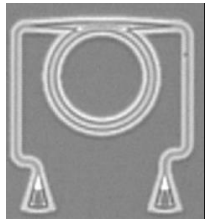

Source

Delay

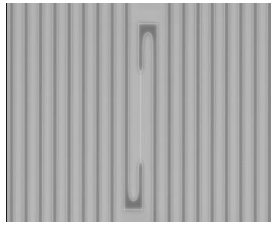

Detector

Source: PsiQuantum

Waveguide propagation loss

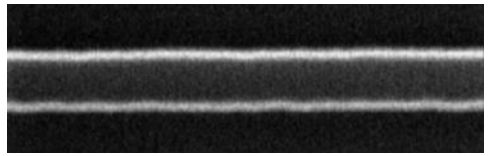

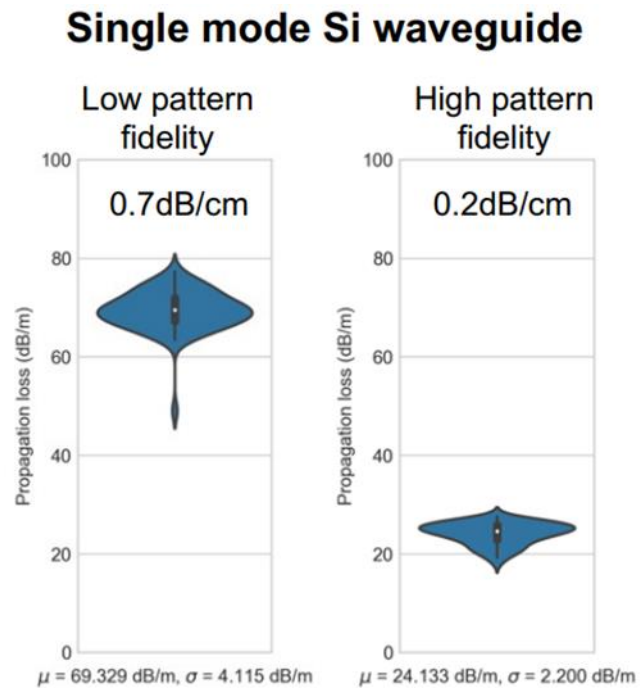

Foundry support for creating quantum computers

- Foundry platform technologies address certain portions of the quantum computing solution
  - Examples: 22FDX®, 3D integration, Silicon Photonics
- Using photons as qubits leverages high volume CMOS foundry
  - Mirrors, beam splitters and phase shifters move photons along circuits in chips networked by fiber. Results are read using photon detectors.
- Known silicon technology makes photonic qubits inherently scalable
- New exotic materials must be carefully considered for predictable performance

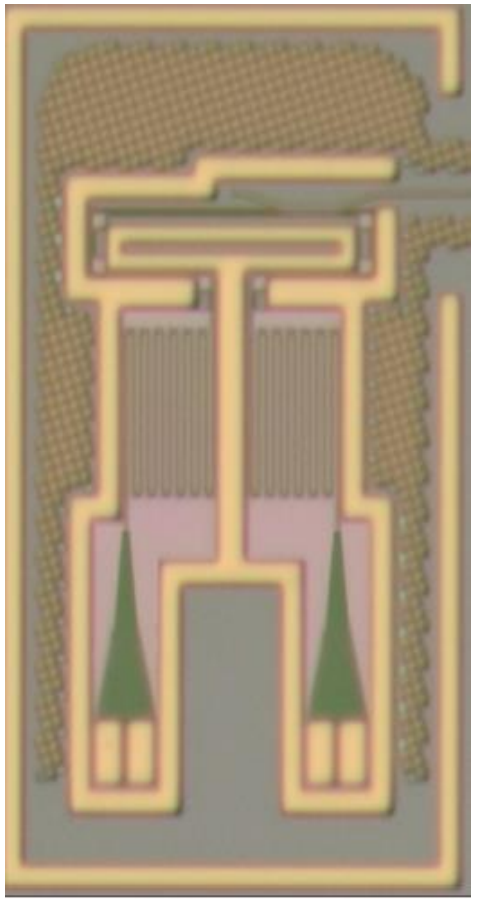

Source: PsiQuantum

Thank You

Premier Sponsor

flexlogix Technologies, Inc.

Platinum Sponsor

ceremorphic®

Gold Sponsors

CEVA
cadence
ARTERIS IP
MOVELLUS
TetraMem
MARVELL
GlobalFoundries
SYNOPSYS® Silicon to Software
Achronix Data Acceleration
proteanTecs
ANDES TECHNOLOGY
expedera
kinara
UNTETHER AI
quadric
innatera
brainchip
Imagination
SIEMENS
BLUE CHEETAH
Qualcomm
NXP
SiFive
intel

TechInsights

Tech Insights

techinsights.com

Thank You!