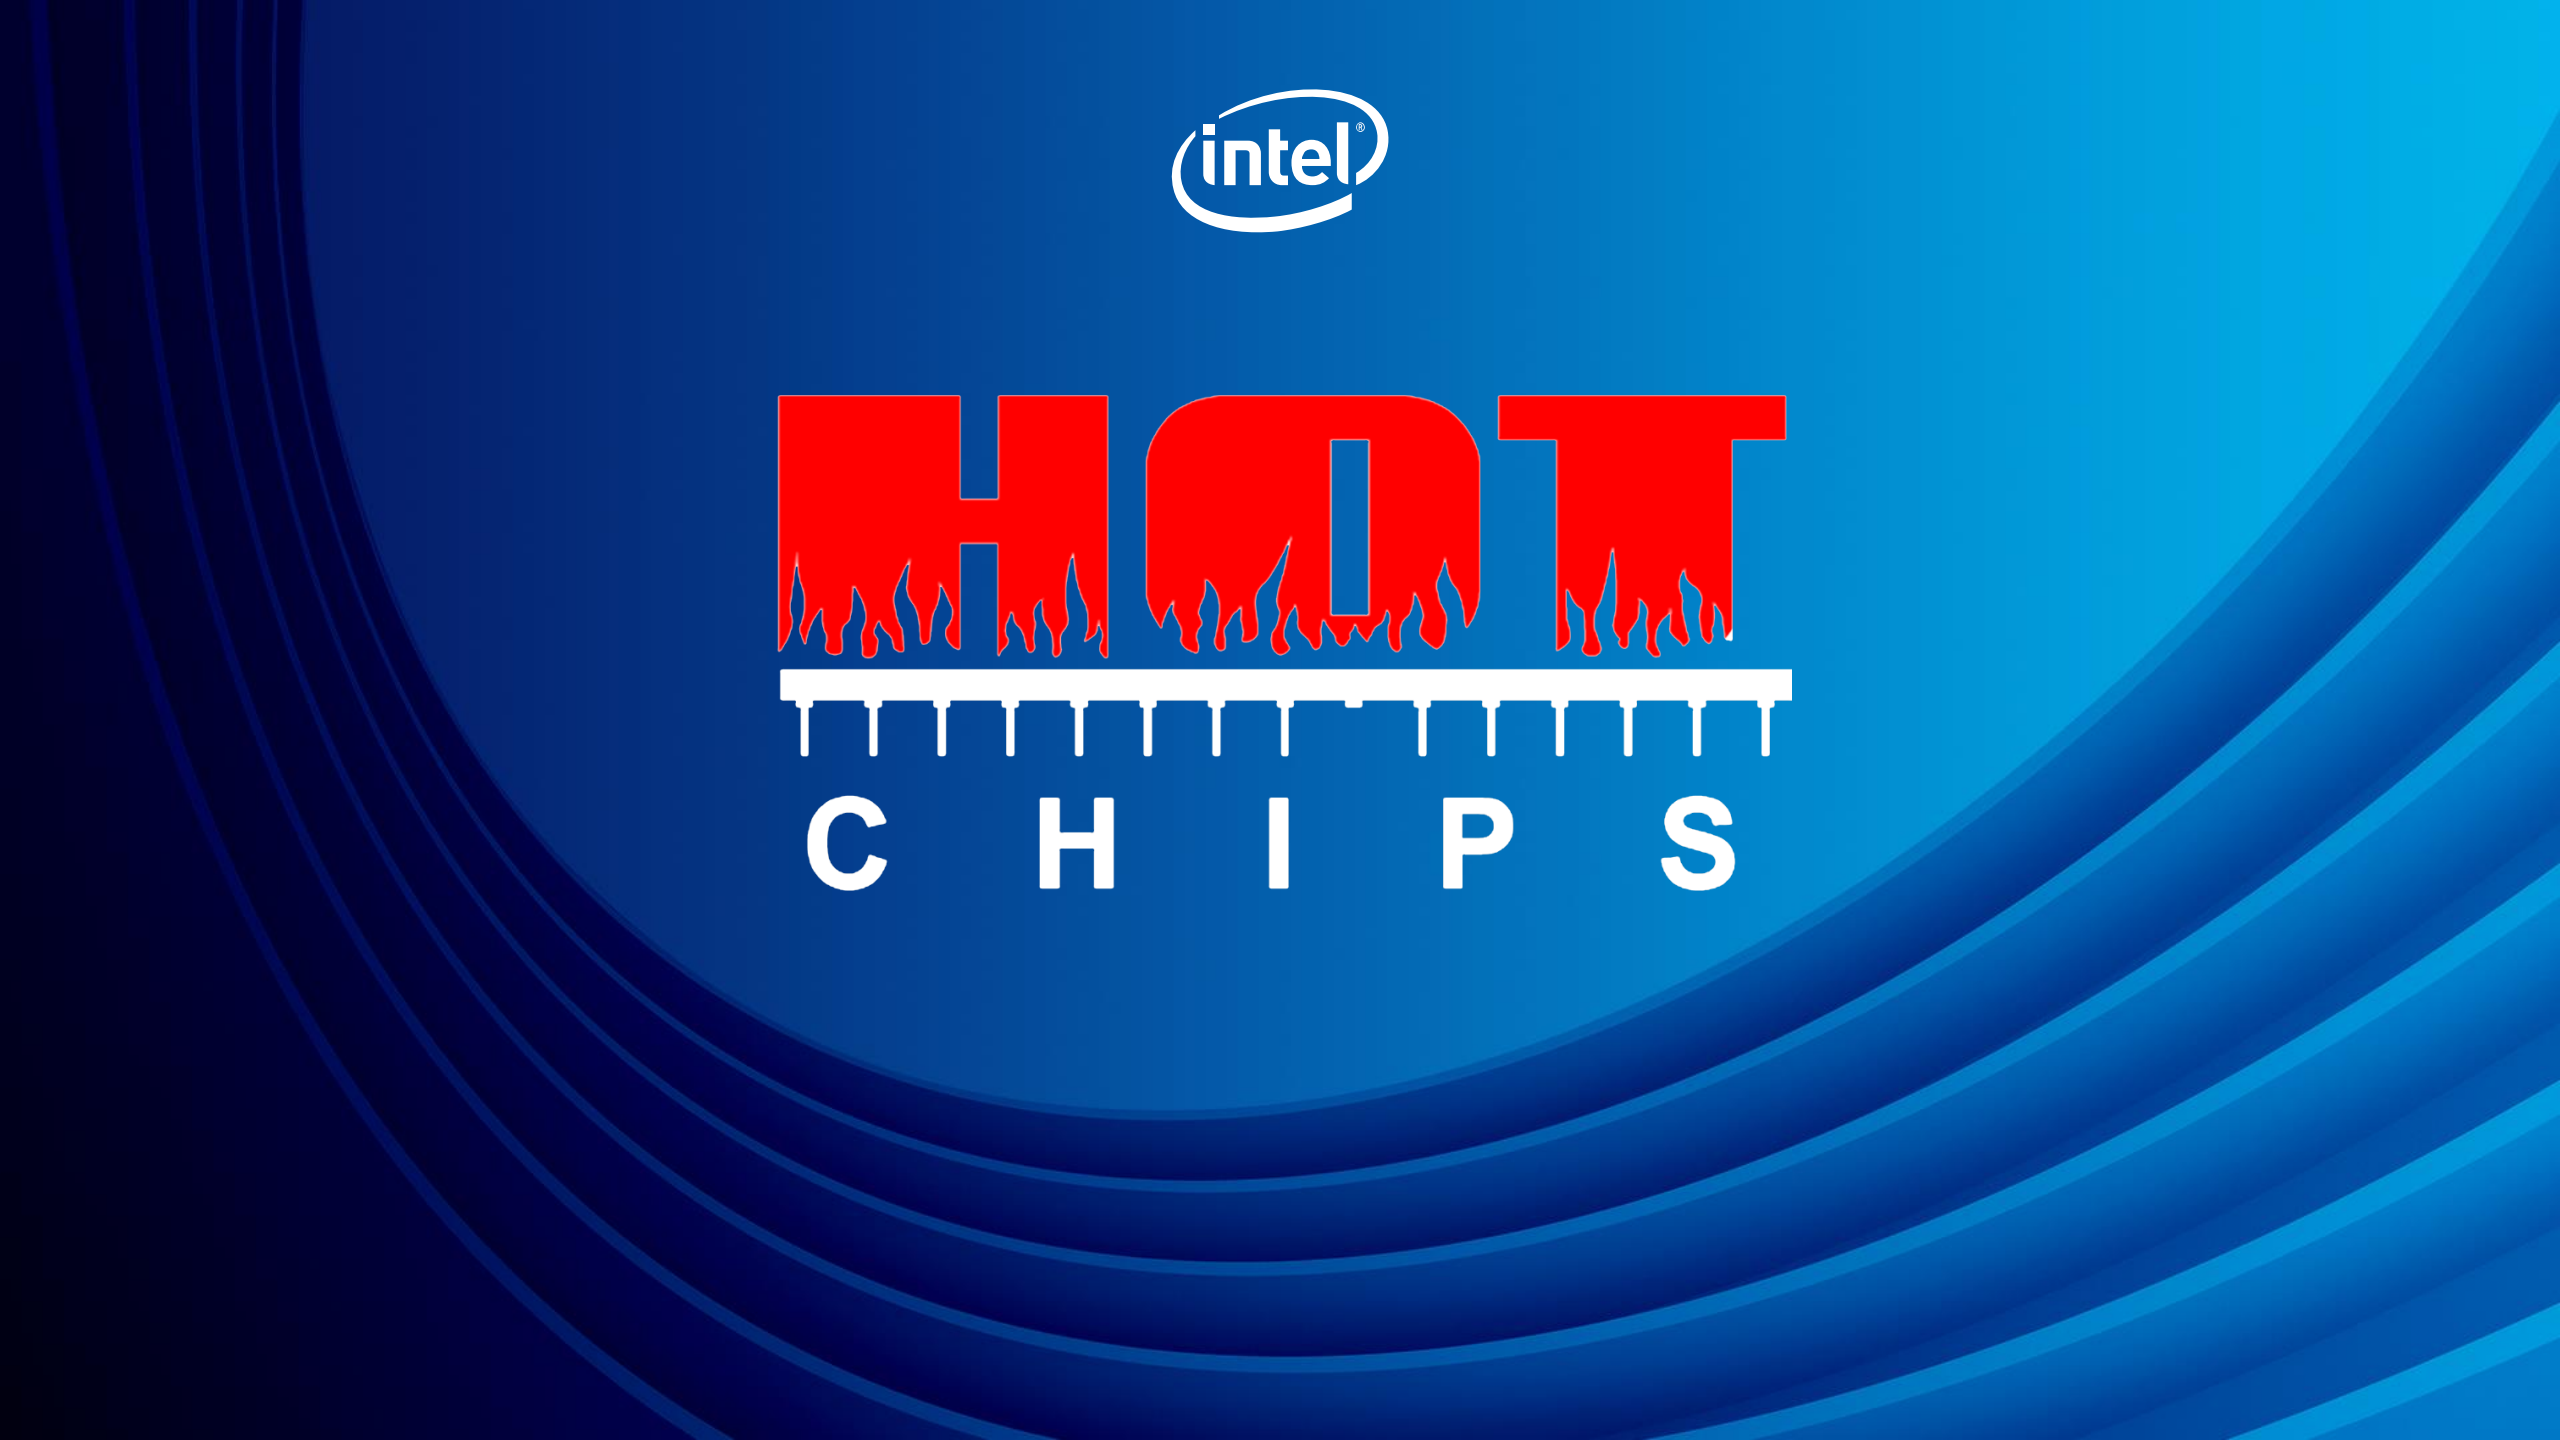
intel
HOT
CHIPS

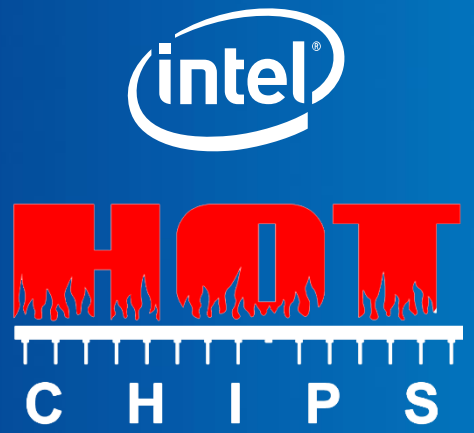

Agilex™ Generation of Intel® FPGAs

Ilya K. Ganusov, Mahesh A. Iyer,

Ning Cheng, Alon Meisler

Intel Corporation

Intel® Agilex™ FPGA for the Data-Centric World

Next-generation core architecture

- New micro-architecture and customized 10 nm process
- 40% higher performance or up to 40% lower power

Advanced network and memory options

- Transceiver tiles with speeds up to 116 Gbps
- Coherent Processor Attach with Compute Express Link (CXL)
- DDR4/5, High-Bandwidth Memory (HBM2e) and Intel Optane™ persistent memory

Extensive software stack

- Leading-edge EDA with Intel® Quartus® Prime software
- OneAPI ecosystem

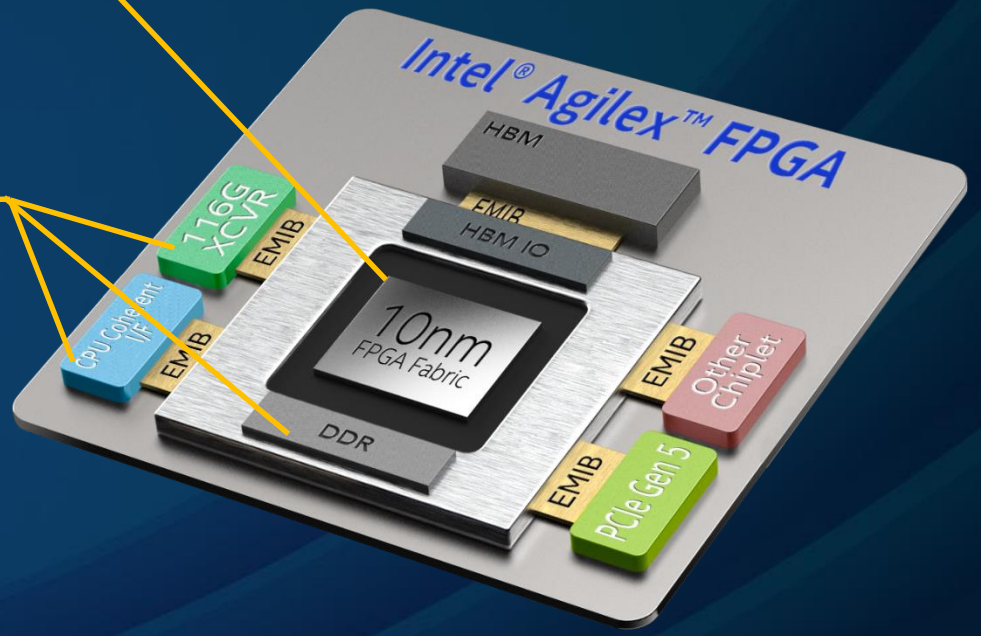

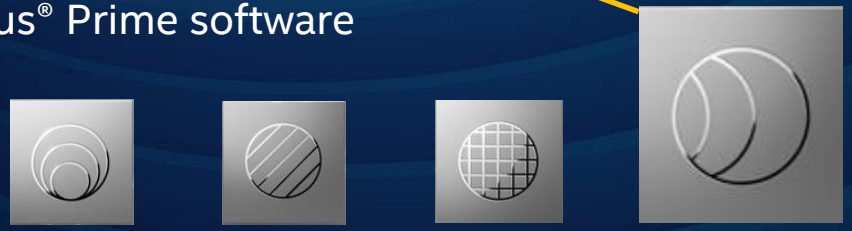

Packaging and Disaggregation

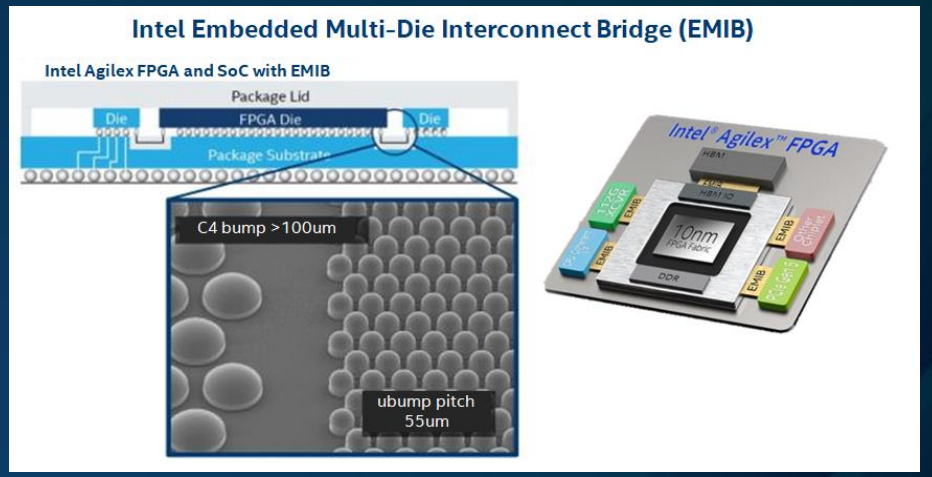

Second-generation EMIB

- High-density interconnect at lower cost compared to Si-interposer
- Flexible chip combination and reuse across different nodes
- Disaggregated transceiver tiles and HBM memory, more coming …

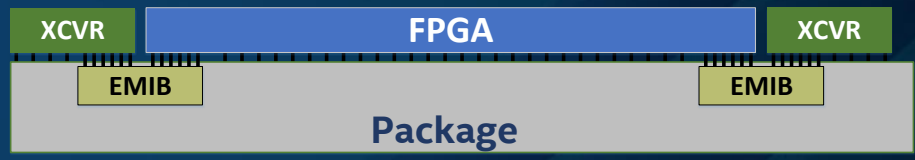

Enables monolithic fabric across full family

- Fabric ease of use while delivering multi-die heterogenous compute

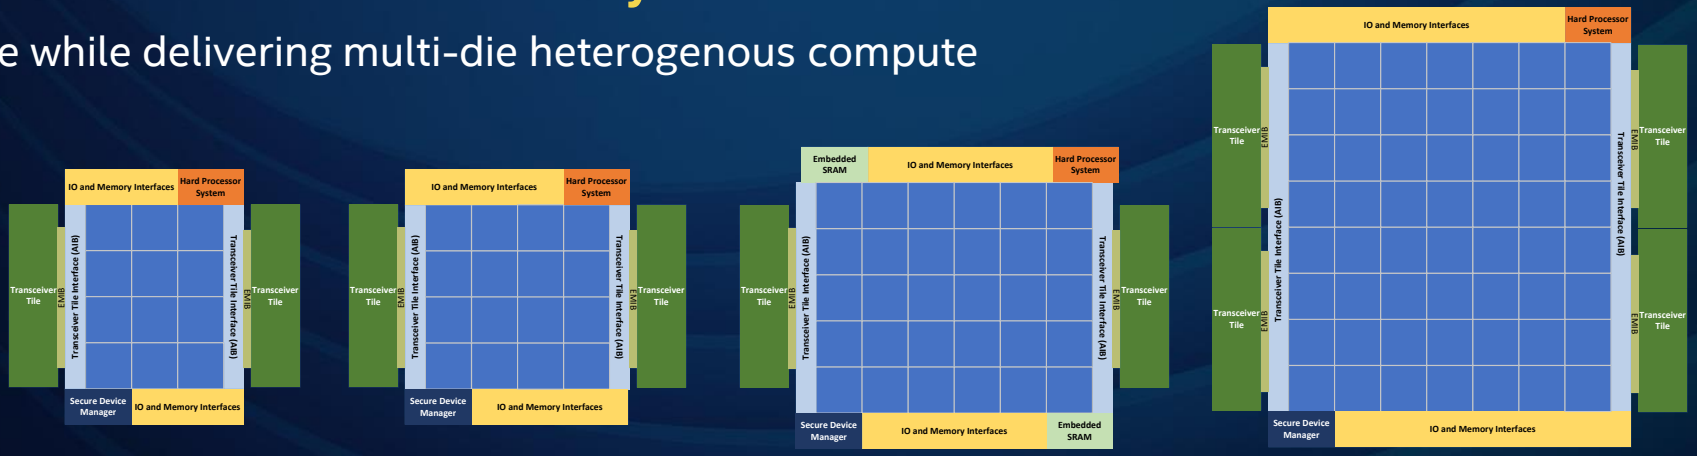

Intel® Agilex™ Device Floorplan

Core Fabric

Smooth fabric grid without I/O column disruptions
No notches in the fabric due to HPS/SDM
Identical resource composition in each sector

Arm64 CPU subsystem

Quad-core A53 with Neon co-processor
ECC L1/L2 caches and snoop control
DMA, USB, UART, SPI, I2C, etc
Coherent I/F between CPU, DDR, and fabric

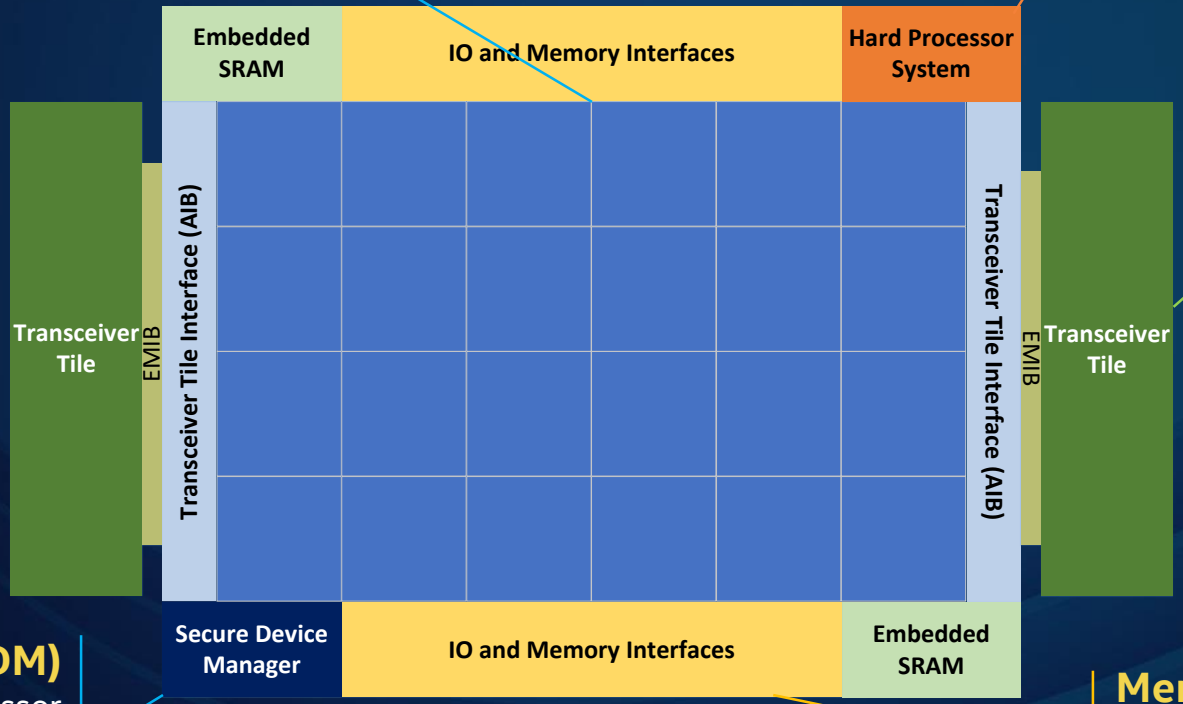

Transceiver Chiplets

16G / 28G / 32G / 56G / 116G PHY
Hard IP for Ethernet, PCIe* Gen4/Gen5
Coherent CPU attach with CXL

Secure Device Manager (SDM)

Triple-redundant hard processor
Crypto and authentication accelerators
Configuration, tamper-detection, single event upset (SEU) response, boot order

Memory and GPIO

1.2 V GPIO and 1.5 V LVDS
DDR/LPDDR and with integrated mem controllers
Intel® Optane™ persistent memory up to 8TB
HBM2e

intel

HOT CHIPS

Intel® 10 nm Process Technology

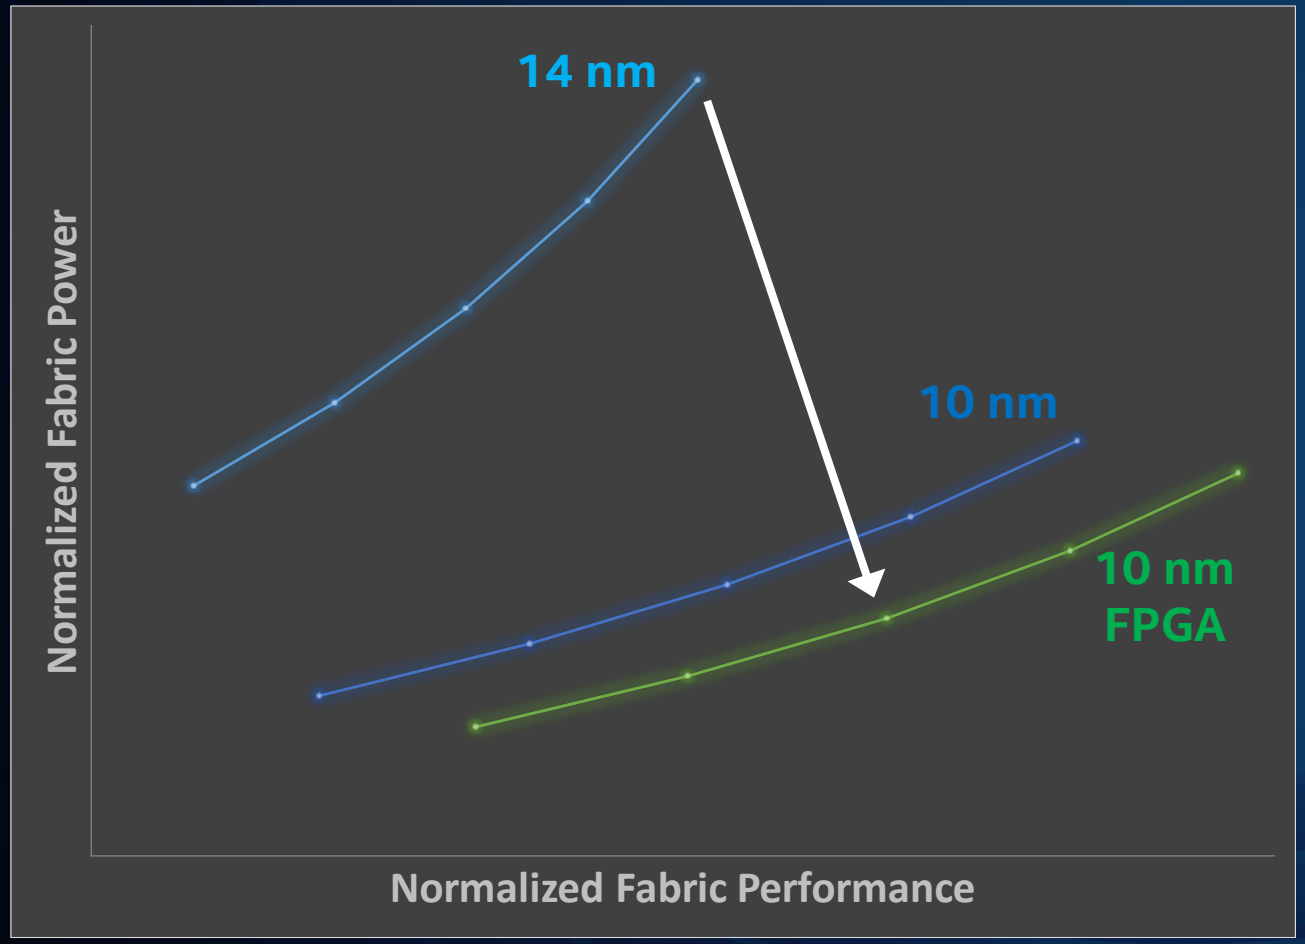

Start with a standard 10 nm logic process

Add extensive customizations for FPGA

- Metal stack and track pattern improvements
- Wider poly pitch for core fabric
- Vt tuning align with Vt optimization
- Fabric custom layout and dummy fill enhancements

Lower Vnom by 100 mV for power reduction

Relying on FPGA core architectural innovation to achieve 40% Fmax increase

Core Fabric Architecture

User-visible resources

- 50% faster Adaptive Logic Module (ALM)
- Extended DSP with 2x throughput for BFLOAT16/FP19/FP16/INT8
- Increased embedded memory block/bit count

Clocking

- Configurable H-tree clocking in sector seams
- Enhanced slack balancing with programmable clock delays

Programmable Routing

- Re-architected high-speed routing
- Enhanced Intel® Hyperflex™ FPGA interconnect registers

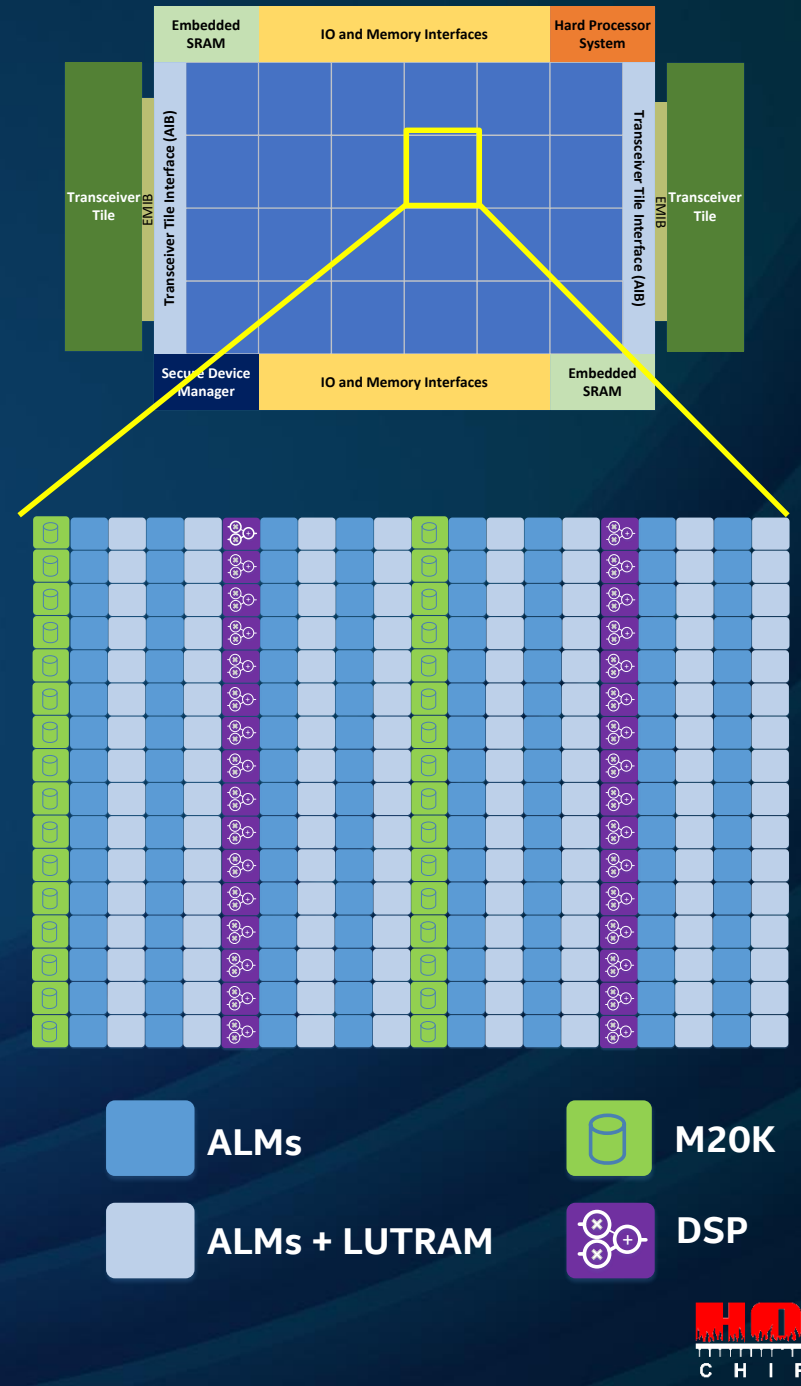

Routing Architecture

Intel® Stratix® 10 FPGA

Wide high-fanout MUXes, multi-drop routing segments

Intel® Agilex™ FPGA

Low-fanout, narrow and fast MUXes, single-drop routing segments

Carefully designed routing pattern to maintain and improve routability

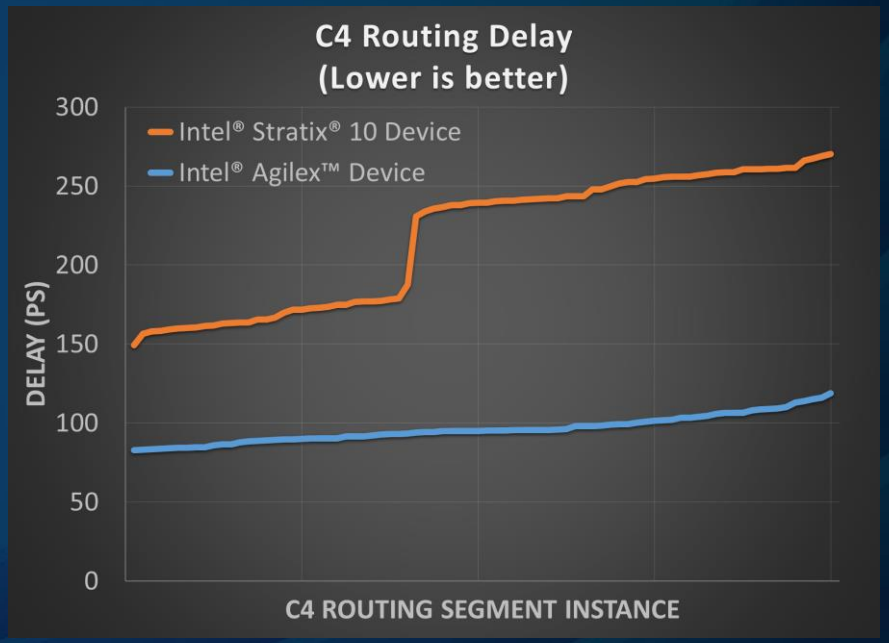

Logic Input Crossbar

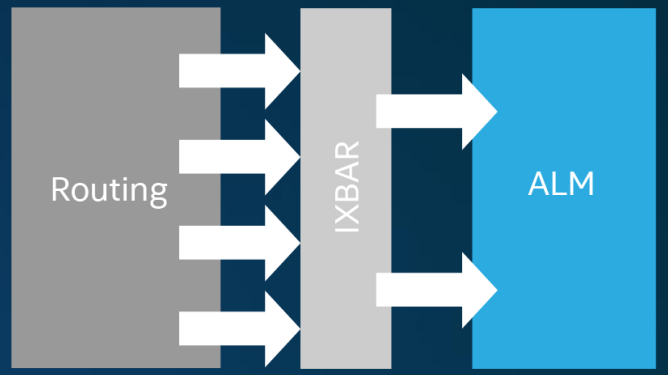

Intel® Stratix® 10 FPGA

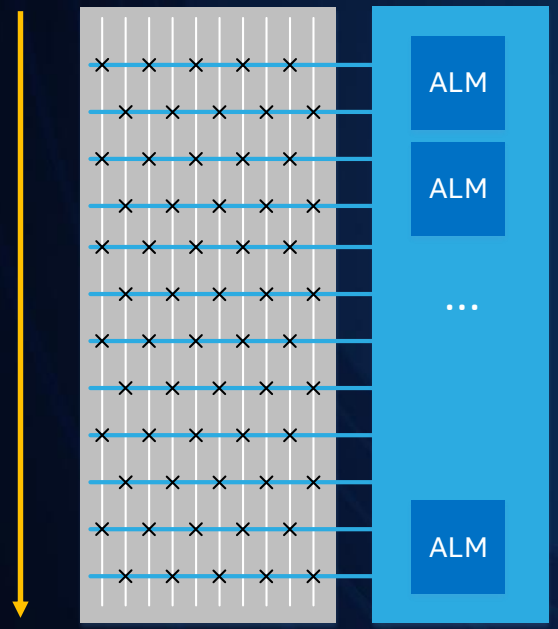

Wide muxes

Intel® Agilex™ FPGA

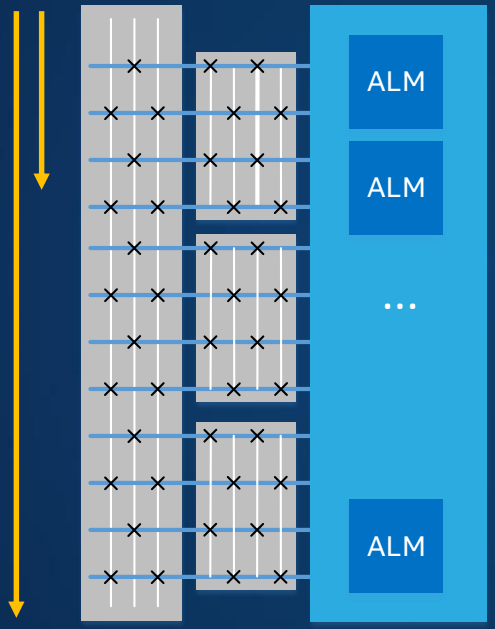

Narrower faster muxes

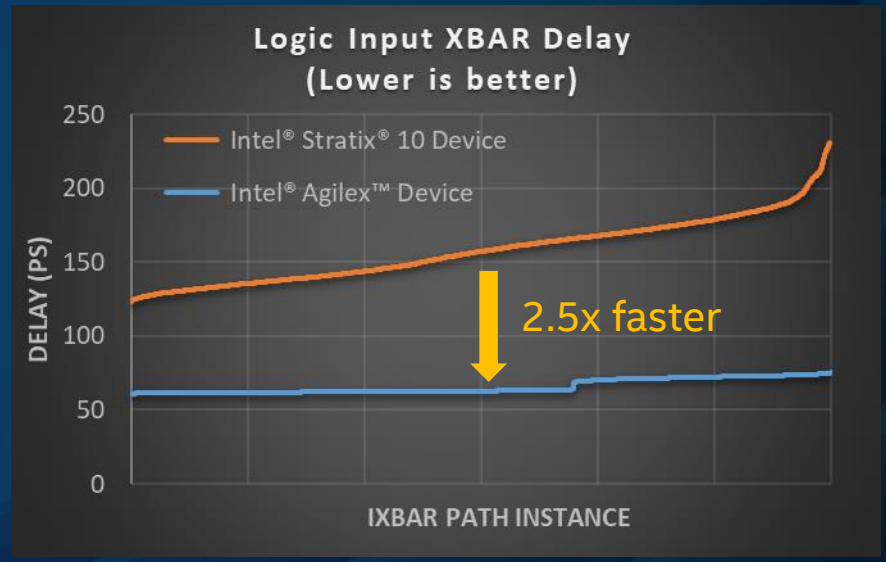

Enhanced CAD
to align logic placement to faster IXBAR lanes

FPGA Fabric Performance Results

Baseline device	Intel® Stratix® 10 device (GX 1100)
Intel® Agilex device	Intel® Agilex™ device (AGF 014)
Speed Grade	-2
Intel® Quartus® Prime software	20.3
Benchmark suite designs	65
Design size	24-427k ALMs

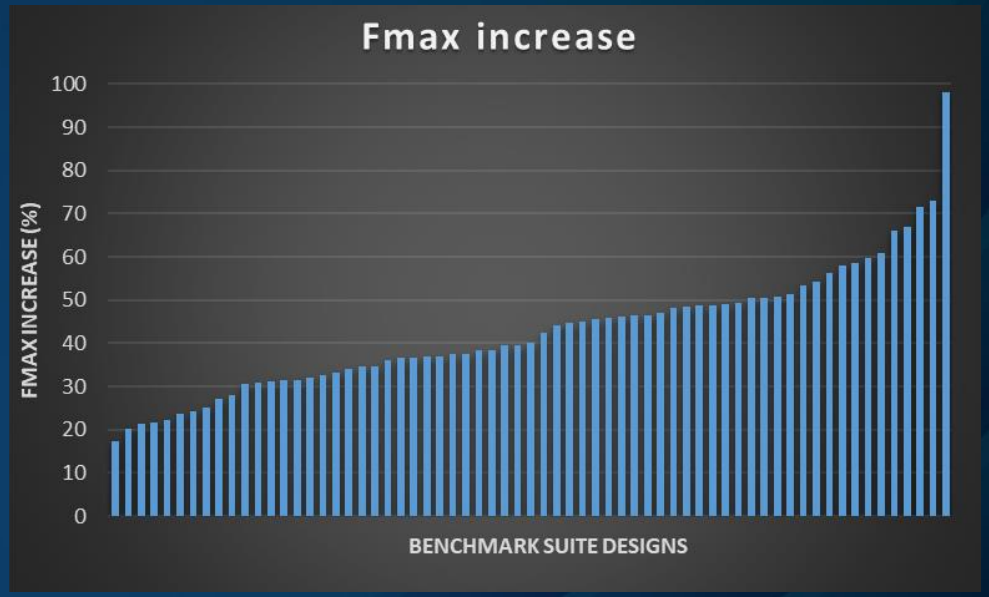

40%+ geomean Fmax improvement over prior generation

- Push-button, same source code
- Broad-based Fmax uplift across entire design suite
- Consistent speedup across wired, wireless, video, ADAS, AI and other application domains

Highest per-generation Fmax uplift since the introduction of Intel® Stratix® FPGAs

Transceiver Tile Choices

PCIe Gen 5 x 16

R-Tile CXL

32G

10/25/50/100 /200/400GE MAC+FEC PCIe Gen4x16

F-Tile

116G

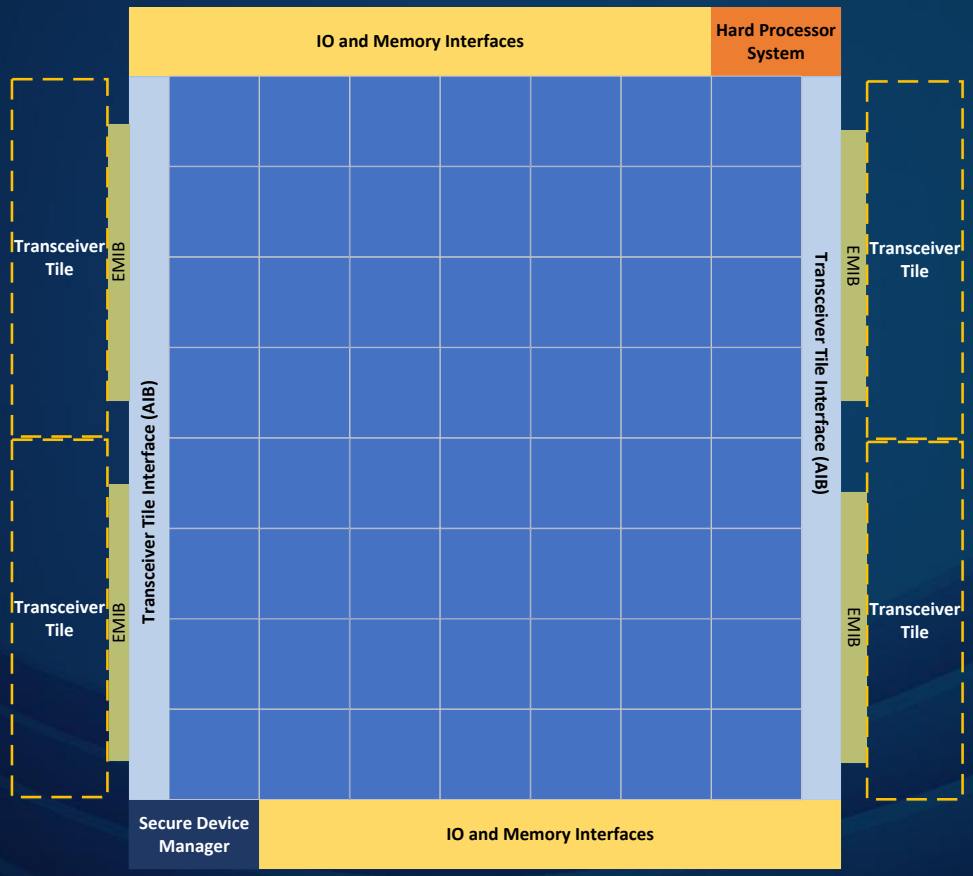

Available in Intel® Agilex™ FPGAs

PCIe Gen 4 x 16

P-Tile

16G

50/100GE MAC

PCIe Gen3x16

H-Tile

28G

10/25/100GE MAC+FEC

E-Tile

58G

Available in Intel® Stratix® 10 and Intel® Agilex® FPGAs

R-Tile with PCIe* Gen5 and CXL

Silicon-proven transceiver PHY meeting Gen5 specification

- 16 lanes of 32 Gbps non-return to zero (NRZ) per tile
- Configurable pipe-direct bypass to FPGA fabric

PCIe* Gen5/CXL common capabilities

- PCIe 1x16, 2x8, and 4x4 RP
- CXL 1x16, 1x8
- SR-IOV 8PF/2kVF PCIe/256VF CXL, VirtIO on PF/VF, scalable IOV
- Config bypass for customized PCIe reg topologies
- Precise time management

Compute Express Link (CXL) coherency attach

- First CXL-attached device with full support for Type 1/2/3 configuration
- Coherency opening new use cases in heterogeneous CPU/FPGA workloads
- Balance between efficiency of hardened controller logic and flexibility of FPGA fabric

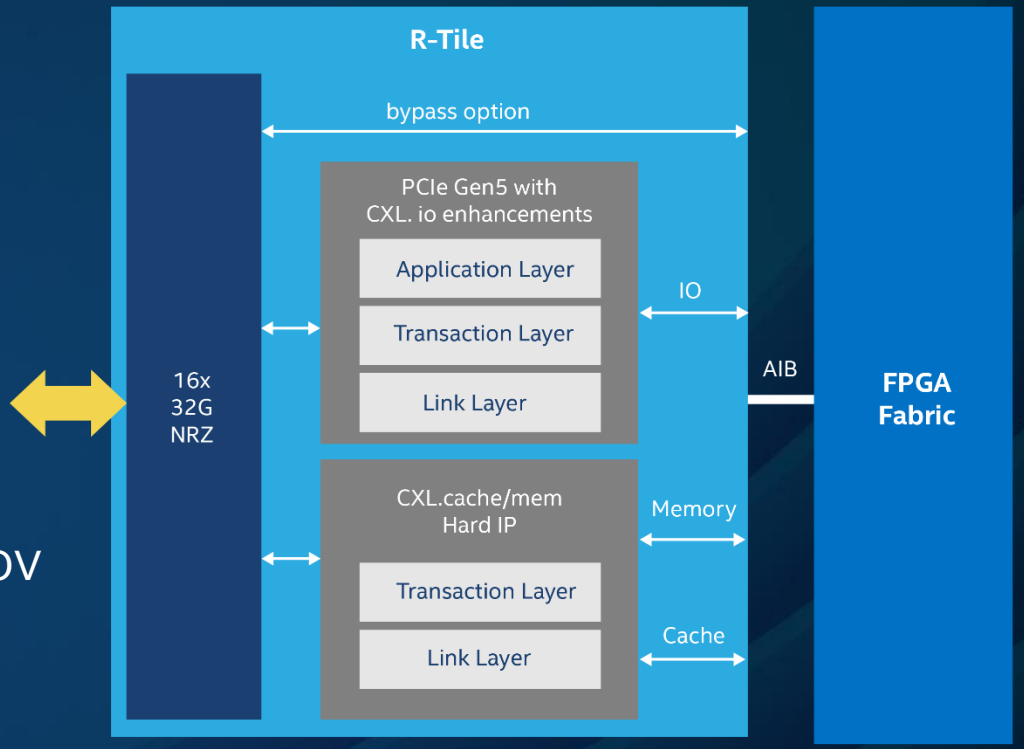

F-Tile with 32G/58G/116G Transceivers

Transceiver PHYs

- 4 x 116G PAM4
- 12 x 58G PAM4 OR 16 x 32G NRZ

Advanced networking with hard IP for MAC, PCS, FEC

- Bifurcatable into combinations 10/25/40/50/100/200/400 GbE
- 600G Interlaken or FlexE
- IEEE 1588 support

Supporting 40+ additional protocols, including

- CEI, CPRI, JESD204B/C, FC, IB, SRIO, GPON, SerialLite FlexO, SDI, SONET/OTN, HDMI, DP, SATA, ...

PCIe Gen4x16

- Precise time management and PMA direct

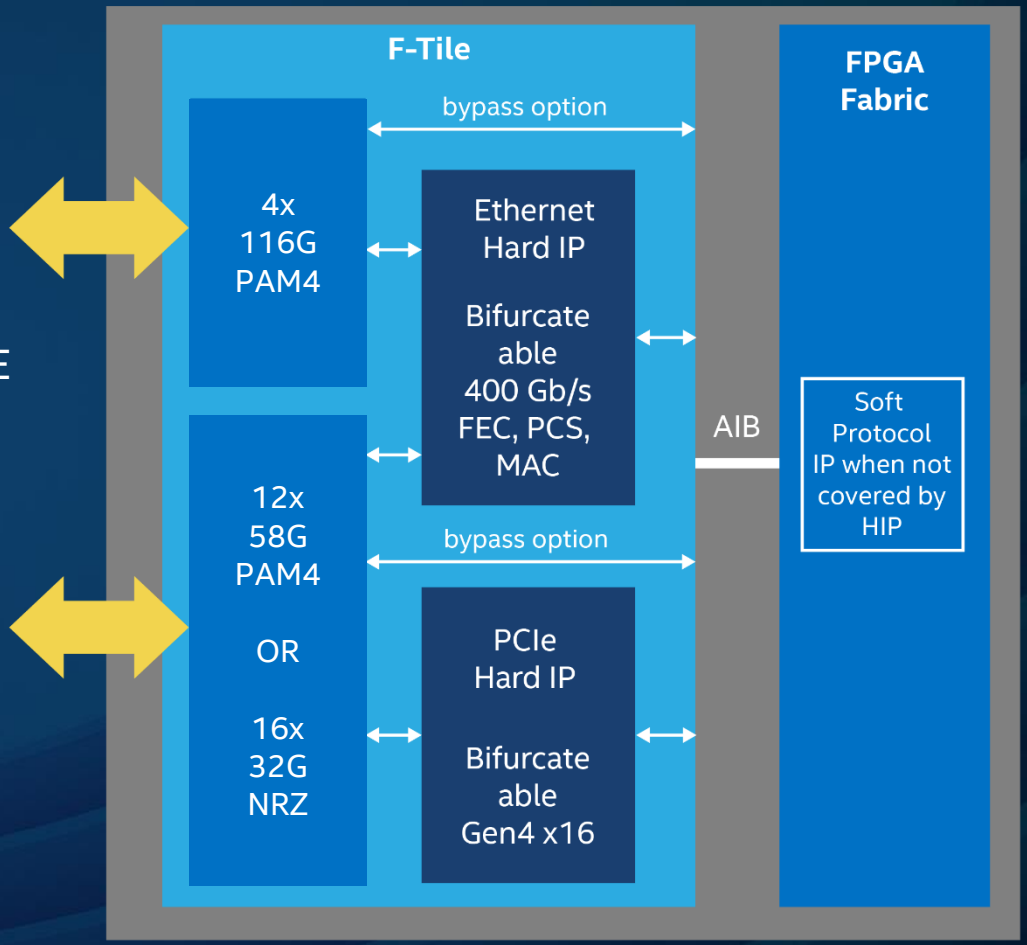

116G Transceivers

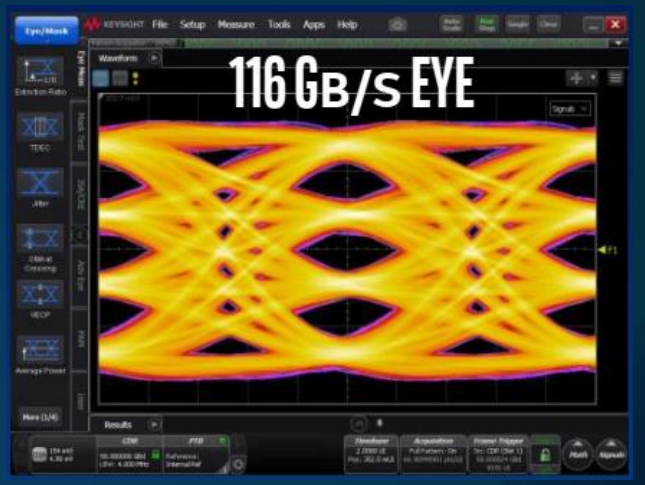

First silicon-proven 116G very long reach with Intel® 10 nm FinFET

- CEI spec compliance across 106G/112G/116G data rates
- >32 dB (ball-to-ball) insertion loss, and ultra-low raw-BER (≈1e-7)
- 116 Gbps rate adds margin and enables emerging standards

Item	Spec.	Unit	Measured Results		
			106.25G	112G	116G
Baud Rate	36-58	GSym/s	53.125	56.0	58
AC Common-Mode Voltage	<30	mV_{rms}	✓ (pass)	✓ (pass)	✓ (pass)
Differential Output Peak-Peak	<1.2	V	✓ (pass)	✓ (pass)	✓ (pass)
Level Mismatch Ratio	>0.95		✓ (pass)	✓ (pass)	✓ (pass)
Steady State Voltage (T_Vf)	400 – 600	mV	✓ (pass)	✓ (pass)	✓ (pass)
Linear Fit Peak (T_Pk)	0.70 × T_Vf	V	✓ (pass)	✓ (pass)	✓ (pass)
SNDR	>32.5	dB	✓ (pass)	✓ (pass)	✓ (pass)
Jrms	<23	mUI_{rms}	✓ (pass)	✓ (pass)	✓ (pass)
J3u	<115	mUI	✓ (pass)	✓ (pass)	✓ (pass)
Even Odd Jitter	<19	mUI_{pp}	✓ (pass)	✓ (pass)	✓ (pass)

116G TX

- Digital-to-analog converter (DAC)-based architecture enabling flexible signaling
- FFE has > 5 taps, exceeding CEI requirements
- Low–power with quarter-rate clocking

116 RX

- Analog-to-digital converter (ADC)-based architecture enabling signal processing or correction in digital domain
- ITOL exceeds CEI-106/112/116 Gbps long reach requirements

Intel® Agilex™ Device Family

	F-Series Wide range of applications	I-Series High-performance CPU attach and bandwidth-intensive apps	M-Series Compute-intensive applications with highest memory bandwidth requirements
Logic Capacity	400k-2.7M LE	2.2 - 2.7M LE	over 2.7M LE
On-chip Memory	5-35 MB	29 - 35 MB	over 35 MB
DSP blocks	up to 8.5k	6 - 8.5k	over 8.5k
Transceiver speeds	up to 58G	up to 116G	up to 116G
PCIe*	PCIe* Gen4	PCIe Gen5	PCIe Gen5
Off-chip Memory	LP/DDR4	LP/DDR4	LP/DDR5 Intel® Optane™ memory
Secure Device Manager	Triple-modular redundant hard processor, encryption and authentication, tamper detection		
Arm* SoC	Quad-core Arm Cortex-A53 up to 1.41 GHz, Neon co-processor, ECC caches, multiple peripherals		
Coherency Option		Compute Express Link to Intel® Xeon® processor	
High-Bandwidth Memory Option			HBM2e

Software Stack

HW/SW Symbiosis in Intel® Agilex™ FPGAs - Next Generation Intel® Hyperflex™ FPGA Architecture

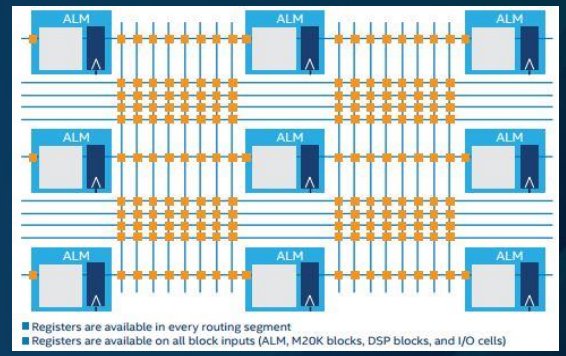

- **Fabric architecture features registers in interconnect**
  - Registers in interconnect routing segments and inputs of logic blocks
  - Fine-grained retiming can break wires without using ALM registers

- **Flexible ASIC-style clocking architecture**

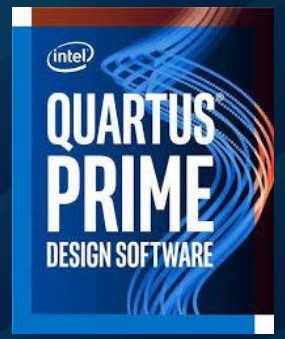

- **Intel® Quartus® Prime Pro Edition software co-designed with FPGA architecture**
  - Fine-grained retiming throughout the flow → Late-stage retiming unique to FPGAs
  - Retiming-awareness and physical clock allocation
  - Fast forward compilation for pipelining recommendations

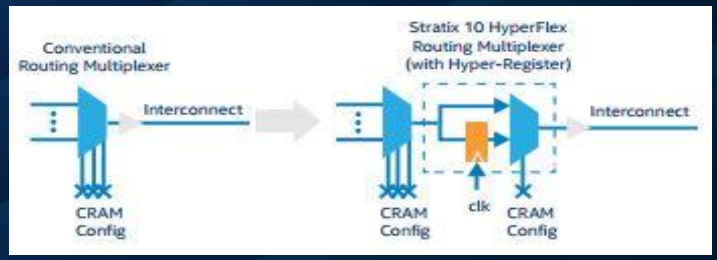

Retiming-Centric FPGA Design Implementation Flow

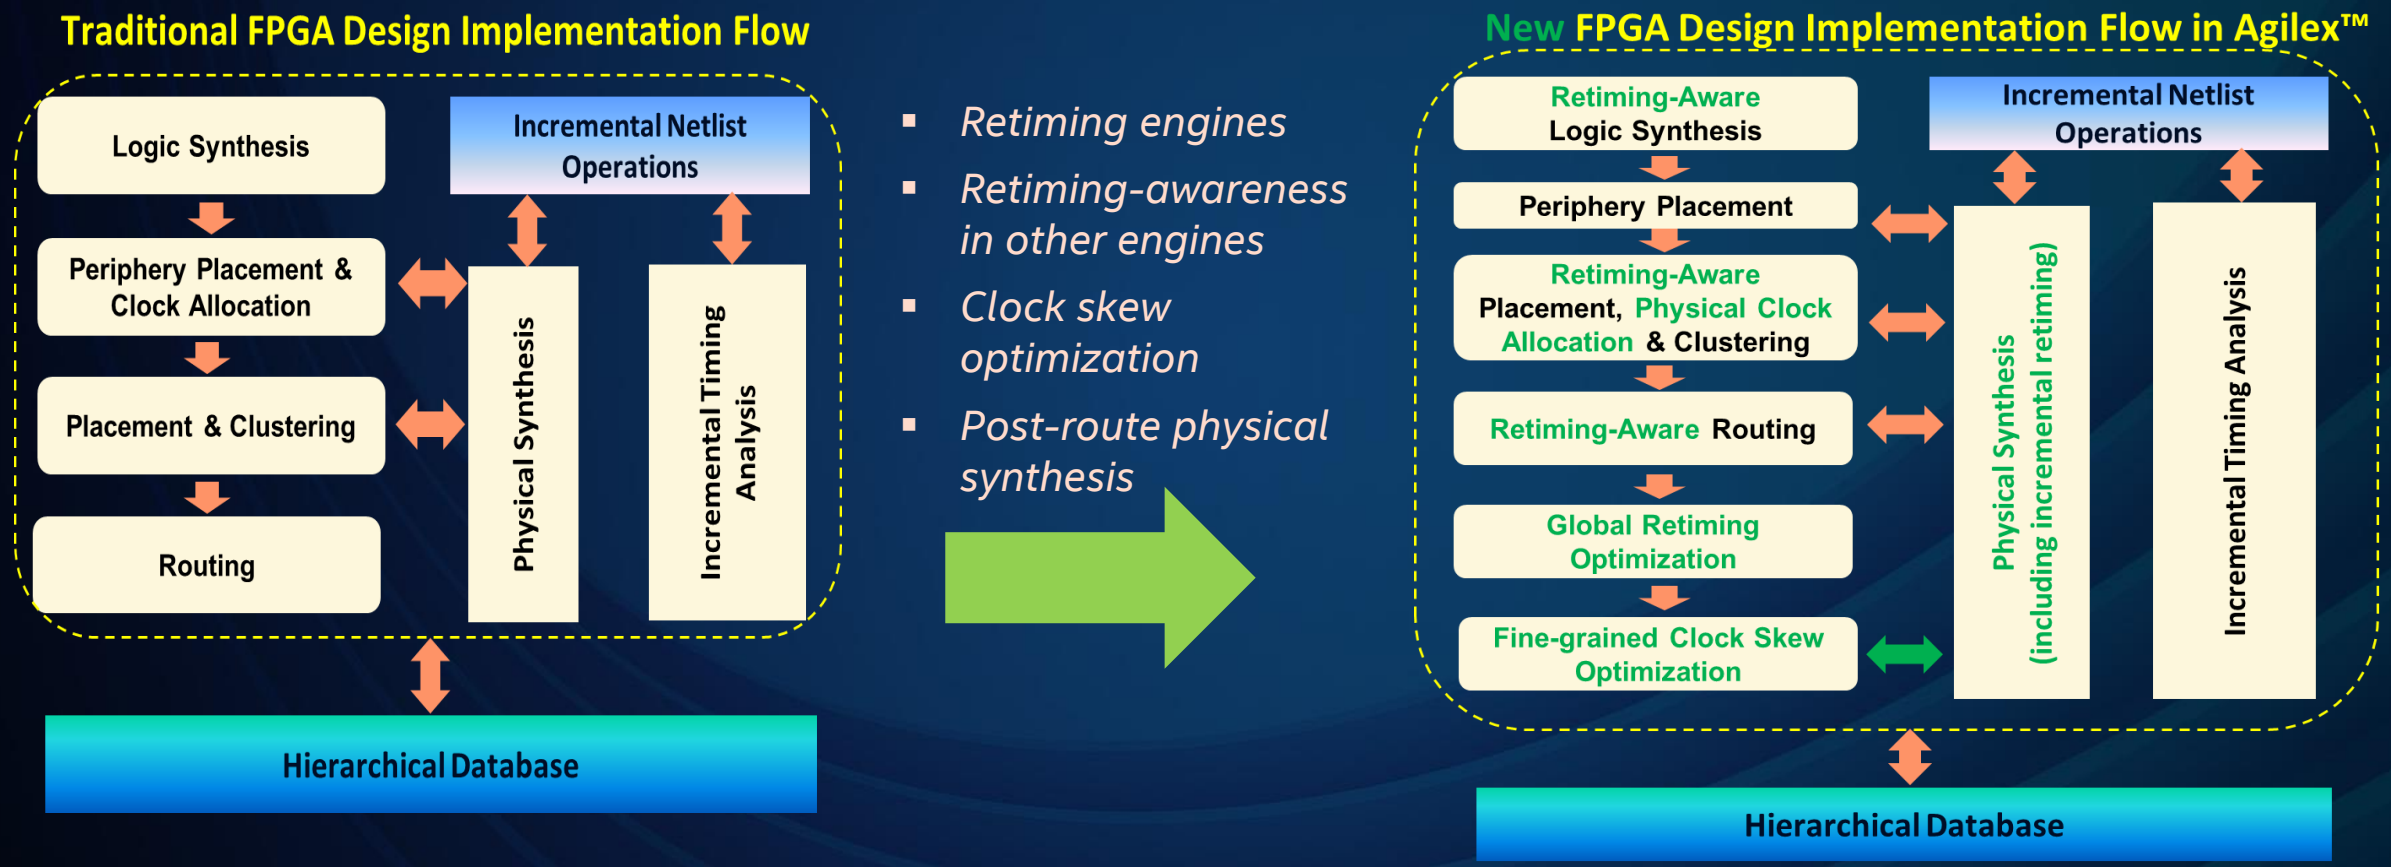

Logic Retiming

- Sequential circuit optimization technique for f_{MAX}
- Optimization objectives: Performance, number of registers
- Retiming modifies circuit structurally

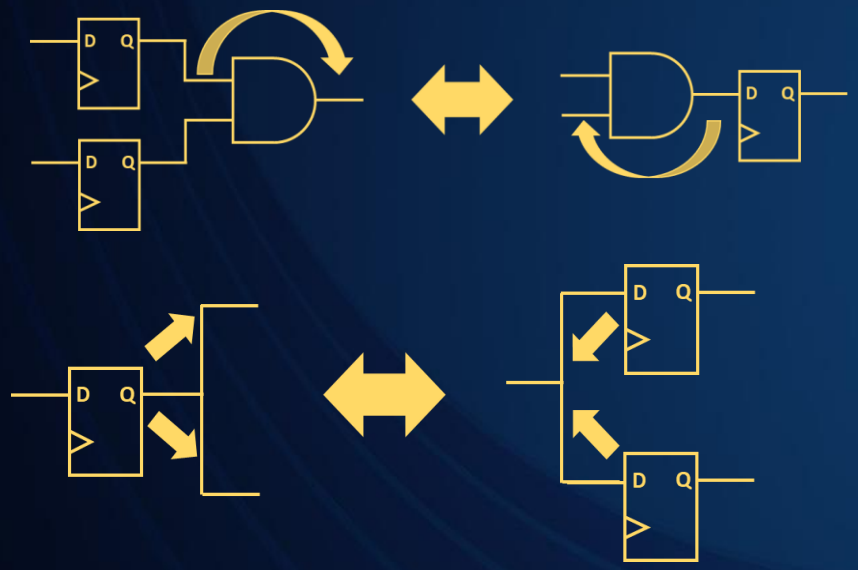

- Graph formulation from Leiserson and Saxe (1991)
  - For an edge from u to v, $\boldsymbol{new\ weight = r_v + old\ weight\ - r_u}$
  - Globally solve with timing constraints

Native Timing-Driven Retiming: U.S. Patent Number: 10,417,374

Sequential Equivalence

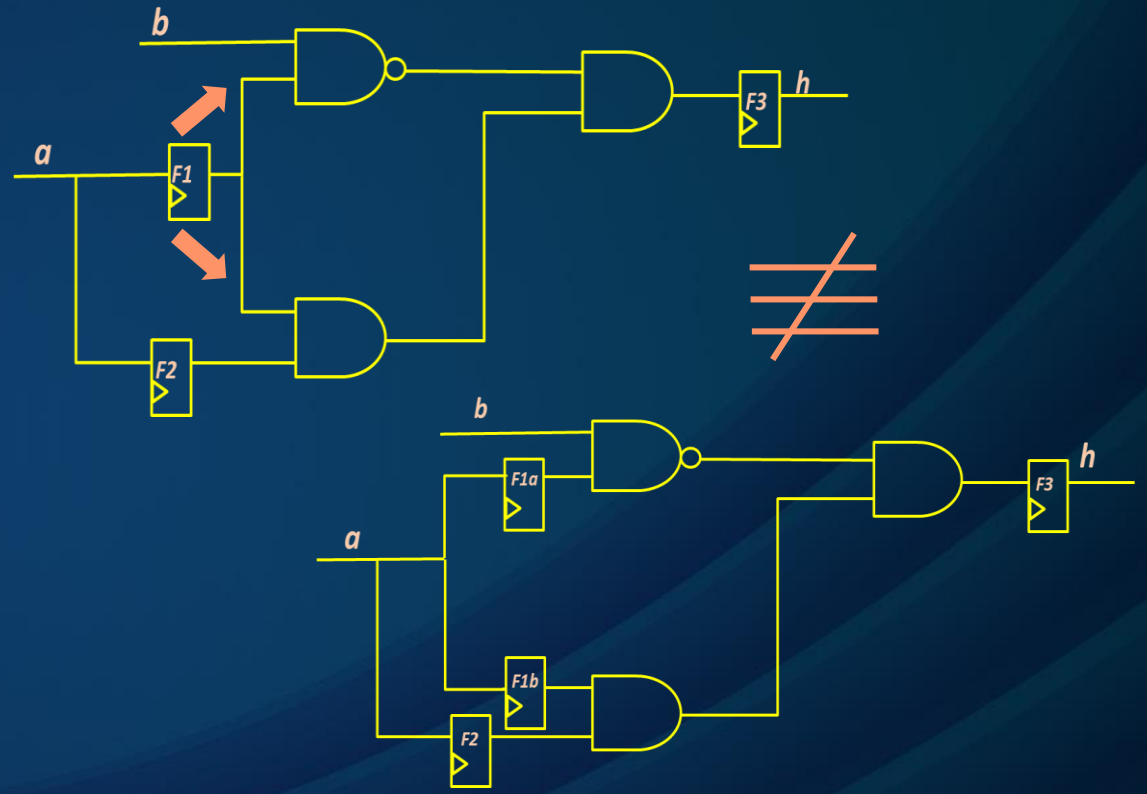

- $\boldsymbol{a = 0, b = 1}$ is a distinguishing sequence
  - For original circuit, this reset sequence produces $\boldsymbol{h = 0}$
  - For initial state, $F1a = 0, F1b = 1, F2 = 1$ this sequence produces $\boldsymbol{h = 1}$

Retiming and Sequential Equivalence

Initial states have to be equivalent in original and retimed circuits

- **Solution 1 → Programmable initial states in HW**
  - Hardware solution costs area
  - Retimer computes initial states

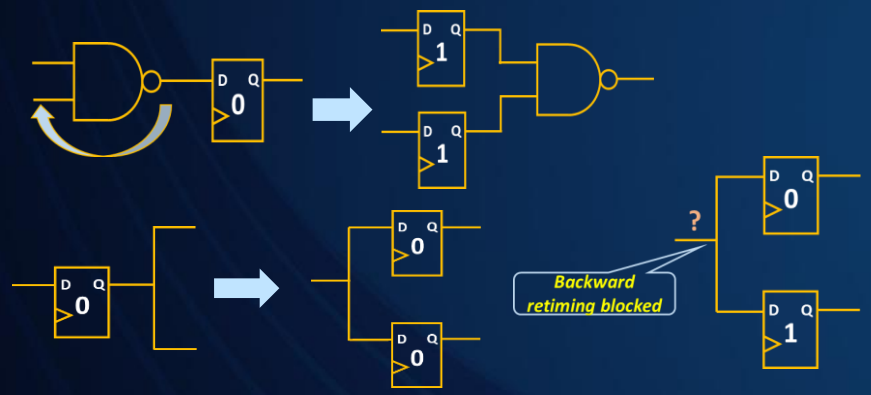

- **Hybrid initial states with don't cares**

U.S. Patent Numbers: 10,296,701, 10,255,404, others pending

- **Solution 2 → *c-cycle* retiming in SW**
  - A retimed circuit is a *c-cycle* delayed replacement of the original circuit
  - Reset sequence for retimed circuit = *c* empty clock cycles + reset sequence of original circuit
  - Retimer computes value of *c* (pessimistic)
  - More flexibility in retiming

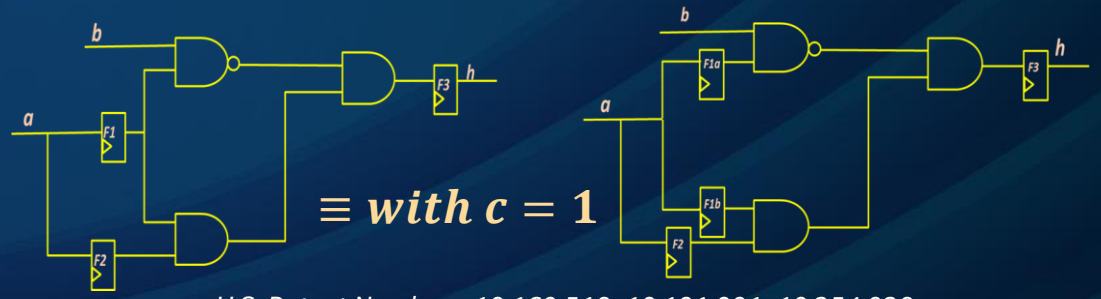

U.S. Patent Numbers: 10,169,518, 10,181,001, 10,354,038

- ***Retiming-Related Optimizations: 25% Average f_{MAX} Increase***
- ***Physical Synthesis Optimizations*: 8% Average f_{MAX} Increase***

** Includes clock skew optimization*

What About Verification of Retimed Circuits?

Retiming has very limited success in ASIC EDA for three decades
LACK OF SCALABLE SEQUENTIAL VERIFICATION TOOLS

- **Breakthrough Rewind Verification of Retimed Circuits**
  - Retiming operations reversible → Attempt to retime retimed circuit back to original circuit
  - For an edge from u to v, $\boldsymbol{new\ weight = r_v + old\ weight\ - r_u}$
  - $\boldsymbol{new\ weight}$ and $\boldsymbol{old\ weight}$ known → solve for $\boldsymbol{r_i}$ variables using constrained random simulation[1]/SAT solver
  - Failure indicates bug in retimer → precise debugging information
  - Initial states verified using reverse-engineered $\boldsymbol{r}$ variable values
- **Significantly faster than a commercial formal verifier**
- **Easily supports c-cycle retiming verification**

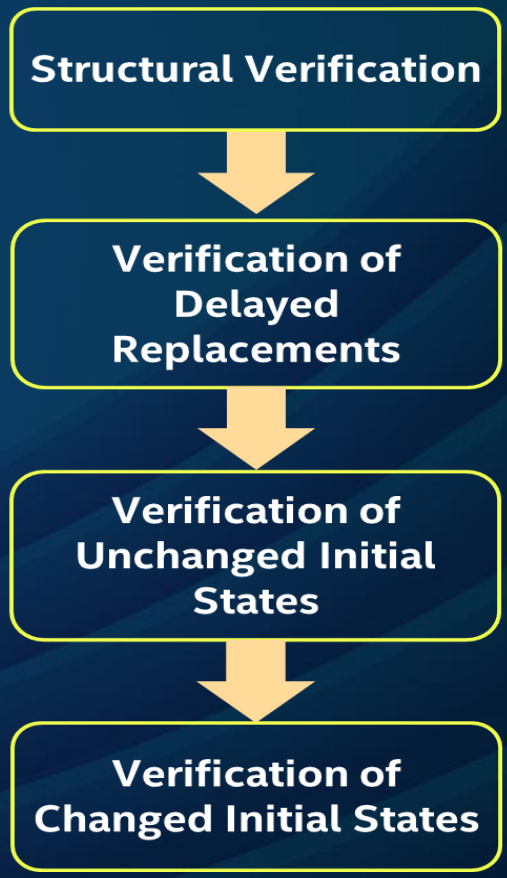

1. *M. A. Iyer, "RACE: A Word-Level ATPG-Based Constraints Solver System for Smart Random Simulation", International Test Conference, 2003*
- *U. S. Patent Numbers: 9,824,177, 10,671,790, 10,157,247, 10,706,203, 10,372,850, 10,489,535, others pending*

Clock Skew Optimization, Time Borrowing, and Retiming

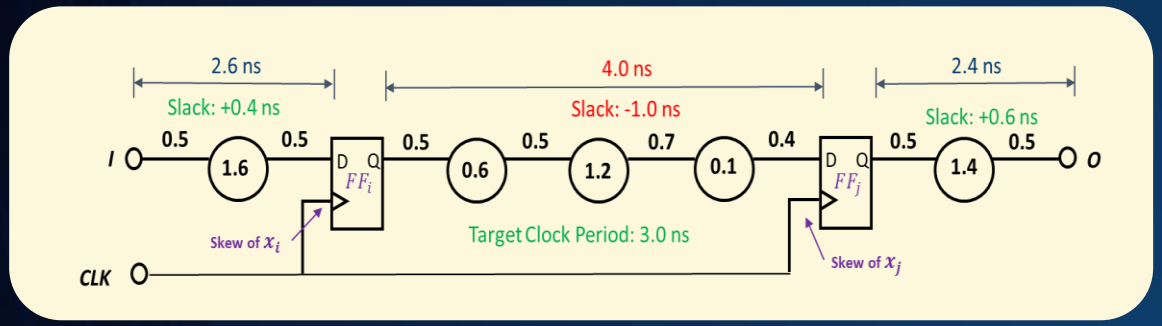

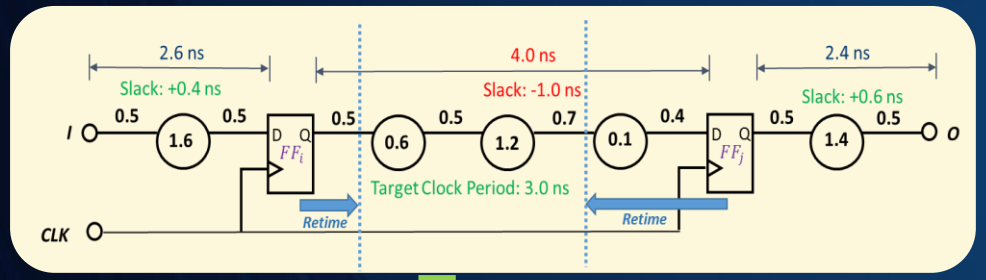

Retiming

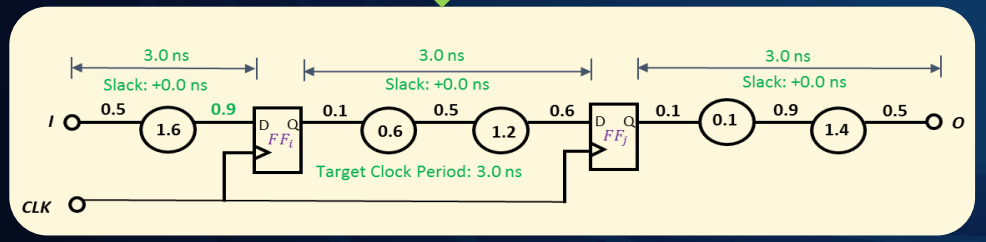

➢ ***Retiming achieves a clock period of 3.0 ns, facilitated using the Intel® Hyperflex™ FPGA Architecture***

- Add clock skew of 0.8 ns on FF_j → Achievable clock period of 3.2 ns
- If FF_i and FF_j are latches, borrow 0.8 ns on FF_j → Achievable clock period of 3.2 ns

f_{MAX} Relationship

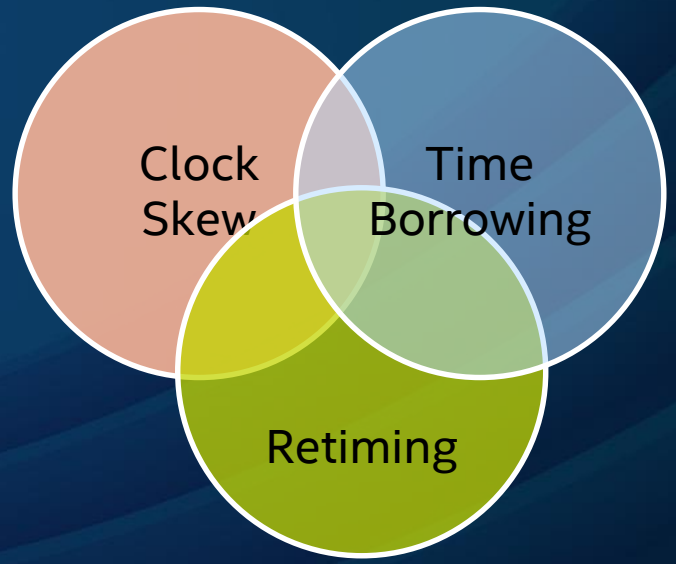

➢ ***Clock Skew Optimization: ~4% average f_{MAX}***

Software-Language Design Entry Options

oneAPI → single unified programming model targeting scalar, vector, matrix, spatial architectures

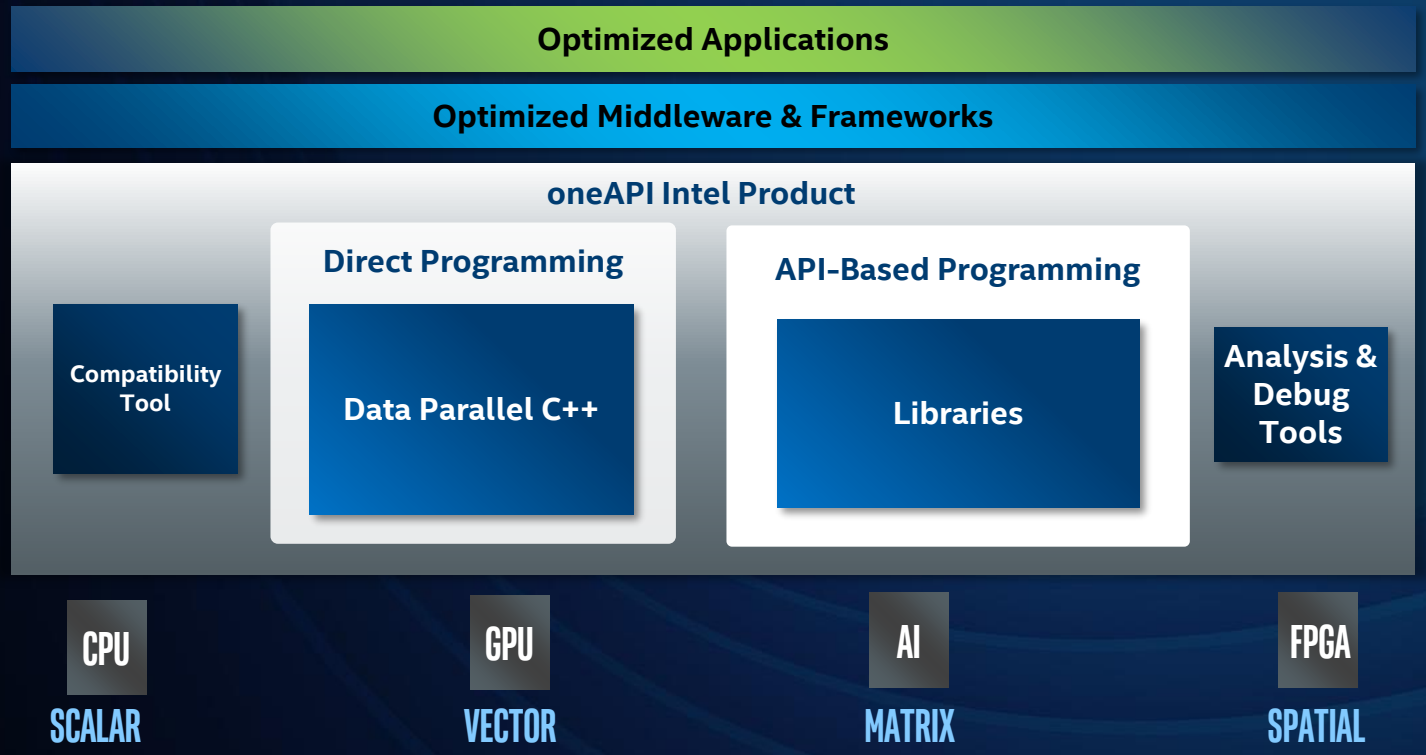

Intel® High-Level Design Tools

- Intel® High Level Synthesis Compiler
- Intel® FPGA SDK for OpenCL™
- DSP Builder for Intel® FPGAs

Notices and Disclaimers

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Performance results are based on testing internal to Intel during Q1 and Q2 of 2020. These tests were done by running internal builds of Intel® Quartus® Prime Pro Design Software on a wide variety of internal benchmarks. The computer systems used for the evaluations were Intel® Xeon® CPU E5-2690 v4 @ 2.60GHz class machines running Red Hat Enterprise Linux Server 6.10 (Santiago) operating system. The performance results represent average improvements across a wide variety of internal benchmarks, and results may vary for each testcase.

Thank You

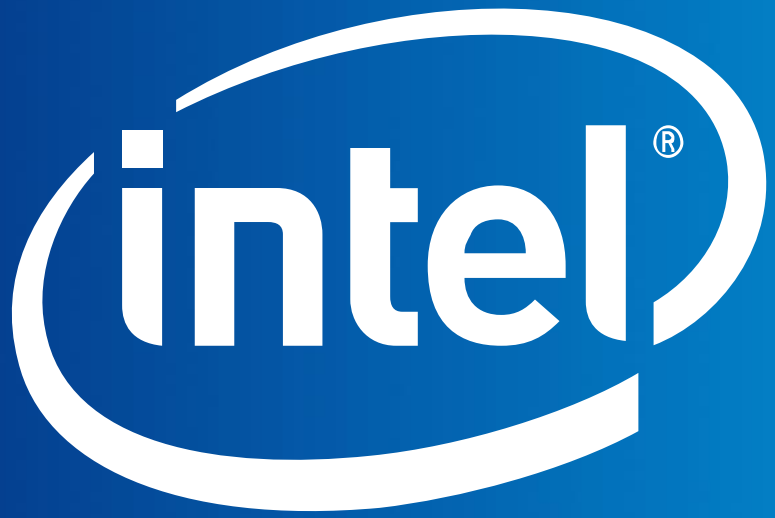
intel