18-643 Lecture 3: FPGA on Moore's Law

James C. Hoe
Department of ECE
Carnegie Mellon University

Housekeeping

- Your goal today: get caught up on 3 decades of progress (upto 2010'ish)
- Notices
  - Complete survey on Canvas, past due
  - Handout #2: lab 0, due noon, 9/11

    Use Piazza and watch TA step-by-step video!!
  - Handout #3: Term Project Intro
- Readings (see lecture schedule online)
  - skim [Boutros, et al., 2021]
  - for next time: skim [Ahmed, et al., 2016] and [Chromczak, et al., 2020]

Where we stopped last time: FPGA as Universal Fabric

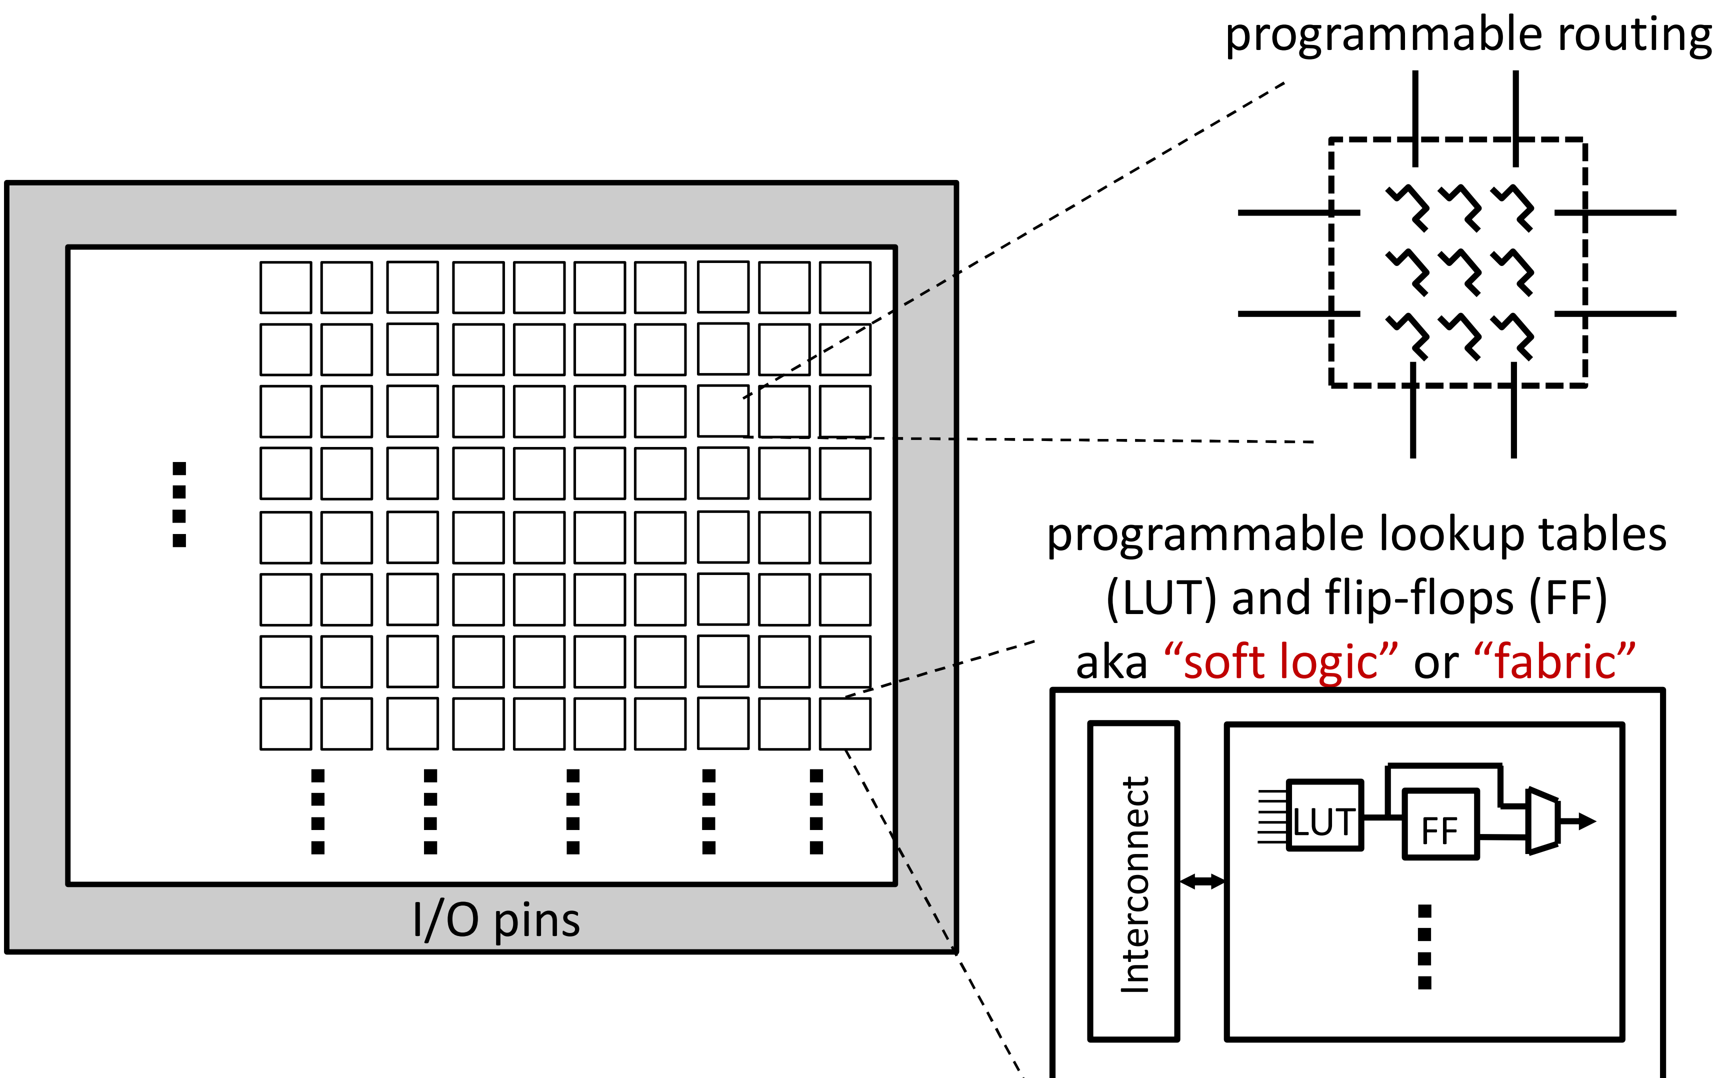

Fast-forward through Moore's Law

Part Number	Logic Capacity (gates)	Configurable Logic Blocks	User I/Os	Configuration Program (bits)
XC2064	1200	64	58	12038
XC2018	1800	100	74	17878

XC2064/XC2018 Logic Cell Arrays: Product Specification

	Kintex UltraScale	Kintex UltraScale+	Virtex UltraScale	Virtex UltraScale+	Zynq UltraScale+
MPSoC Processing System					✓
System Logic Cells (K)	318–1,451	356–1,143	783–5,541	862–3,780	103–1,143
Block Memory (Mb)	12.7–75.9	12.7–34.6	44.3–132.9	23.6–94.5	4.5–34.6
UltraRAM (Mb)		0–36		90–360	0–36
HBM DRAM (GB)				0–8	
DSP (Slices)	768–5,520	1,368–3,528	600–2,880	2,280–12,288	240–3,528
DSP Performance (GMAC/s)	8,180	6,287	4,268	21,897	6,287
Transceivers	12–64	16–76	36–120	32–128	0–72
Max. Transceiver Speed (Gb/s)	16.3	32.75	30.5	32.75	32.75
Max. Serial Bandwidth (full duplex) (Gb/s)	2,086	3,268	5,616	8,384	3,268
Integrated Blocks for PCIe®	1–6	0–5	2–6	2–6	0–5
Memory Interface Performance (Mb/s)	2,400	2,666	2,400	2,666	2,666
I/O Pins	312–832	280–668	338–1,456	208–832	82–668
I/O Voltage (V)	1.0–3.3	1.0–3.3	1.0–3.3	1.0–1.8	1.0–3.3

what happened is more than Moore

[Table 1, UltraScale Architecture and Product Datasheet: Overview]

30 Years of Becoming Hardwired

LUT-based Configurable Logic Block (simplified sketch)

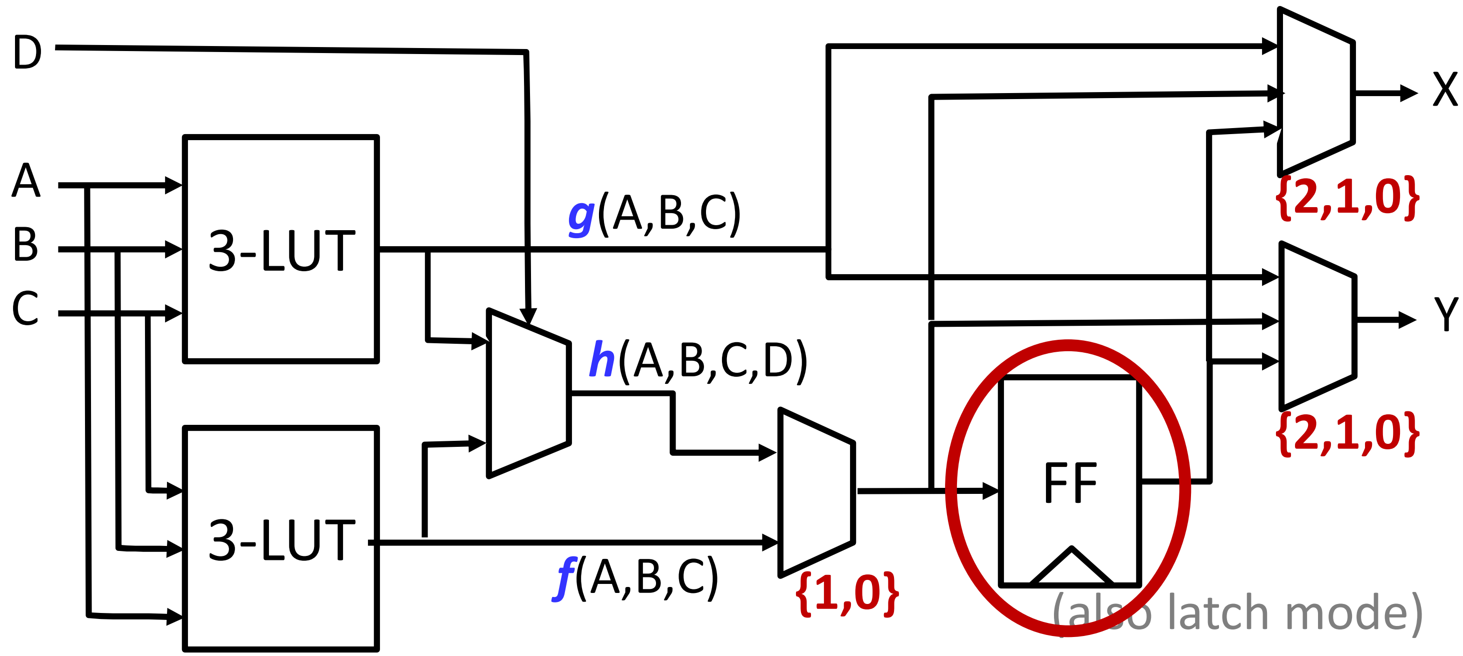

- 2 fxns (***f*** & ***g***) of 3 inputs OR 1 fxn (***h***) of 4 inputs
- hardwired FFs (too expensive/slow to fake)
- Just 10s of these in the earliest FPGAs

Recall

Why Hardwired Logic

- LUTs already can do everything (digital)
- Revisit: why hardwired flip-flop in CLB?
  - would take 4 LUTs to make 1 M-S flip-flop
  - LUT-built FF have atrocious setup/hold time
  - almost all designs affected in cost and speed
- Makes sense to hardwire a functionality
  - needed by everyone (or by the big customers)
  - <u>expected</u> benefit outweigh displaced LUT area, i.e.,
    - much more expensive/slow in LUTs
    - easy/cheap to ignore when not in use

Hardwiring is a great thing if it is usable and is used

E.g., Special Support for Addition

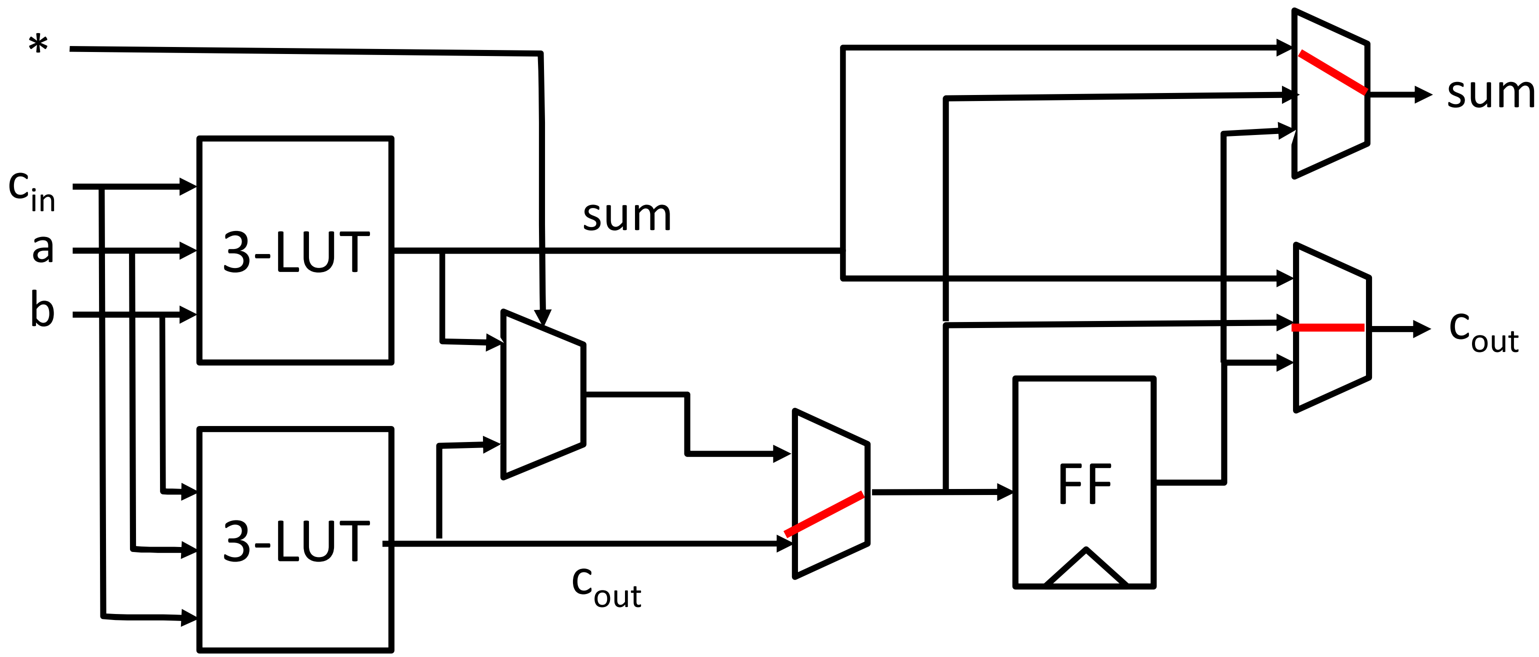

- A full-adder fits perfectly in 1 CLB with 2x3LUTs
- But carry propagation slow---flow through several configurable connections and two switch blocks
- Addition is pretty important to most designs

Specialized Logic for Fast Carry

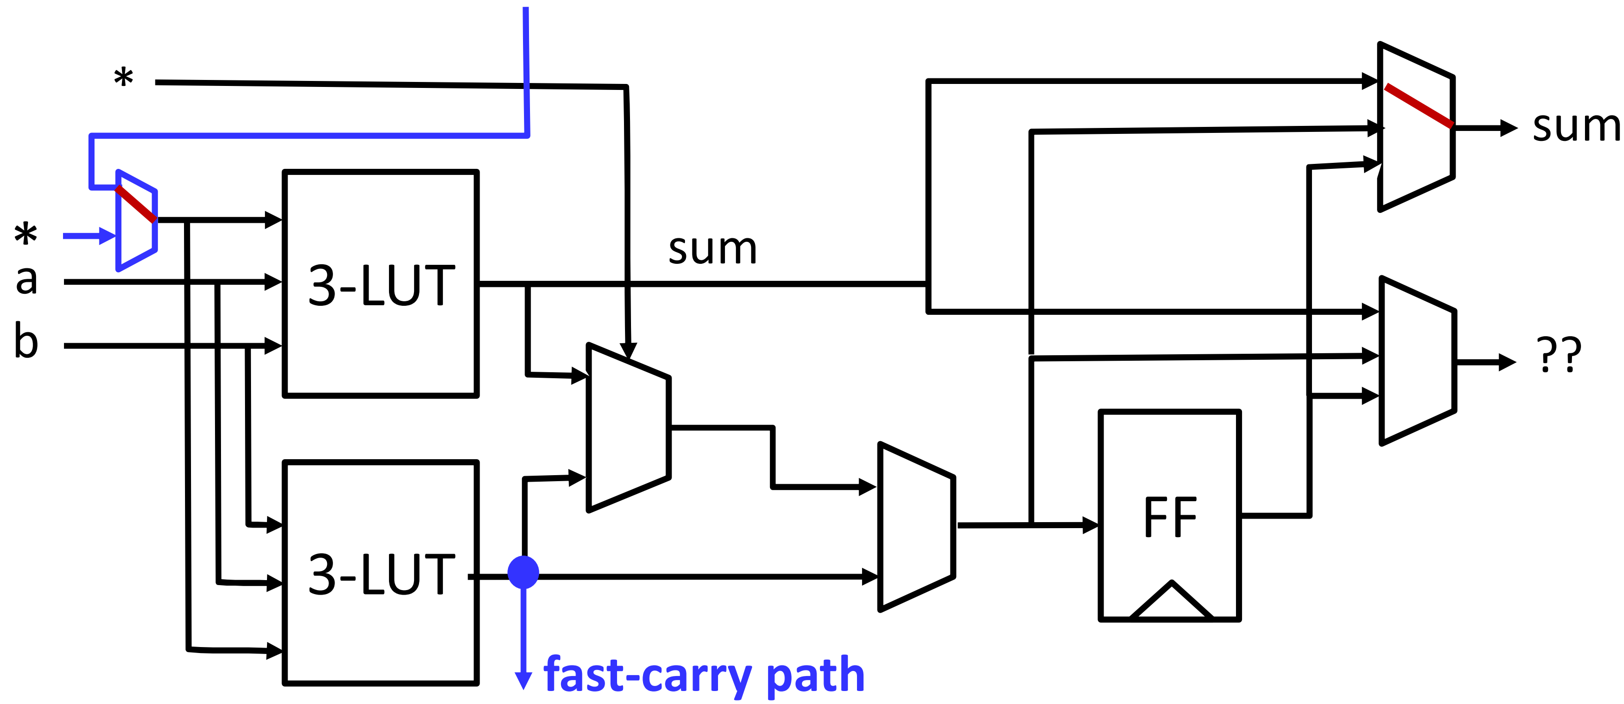

- Cost = 1 (real) wire and 1 mux
- Huge win in adder performance (32-bit@33MHz)

If arithmetic is so important, why not put in real adders? How about multipliers?

Xilinx XC4000 (1990s)

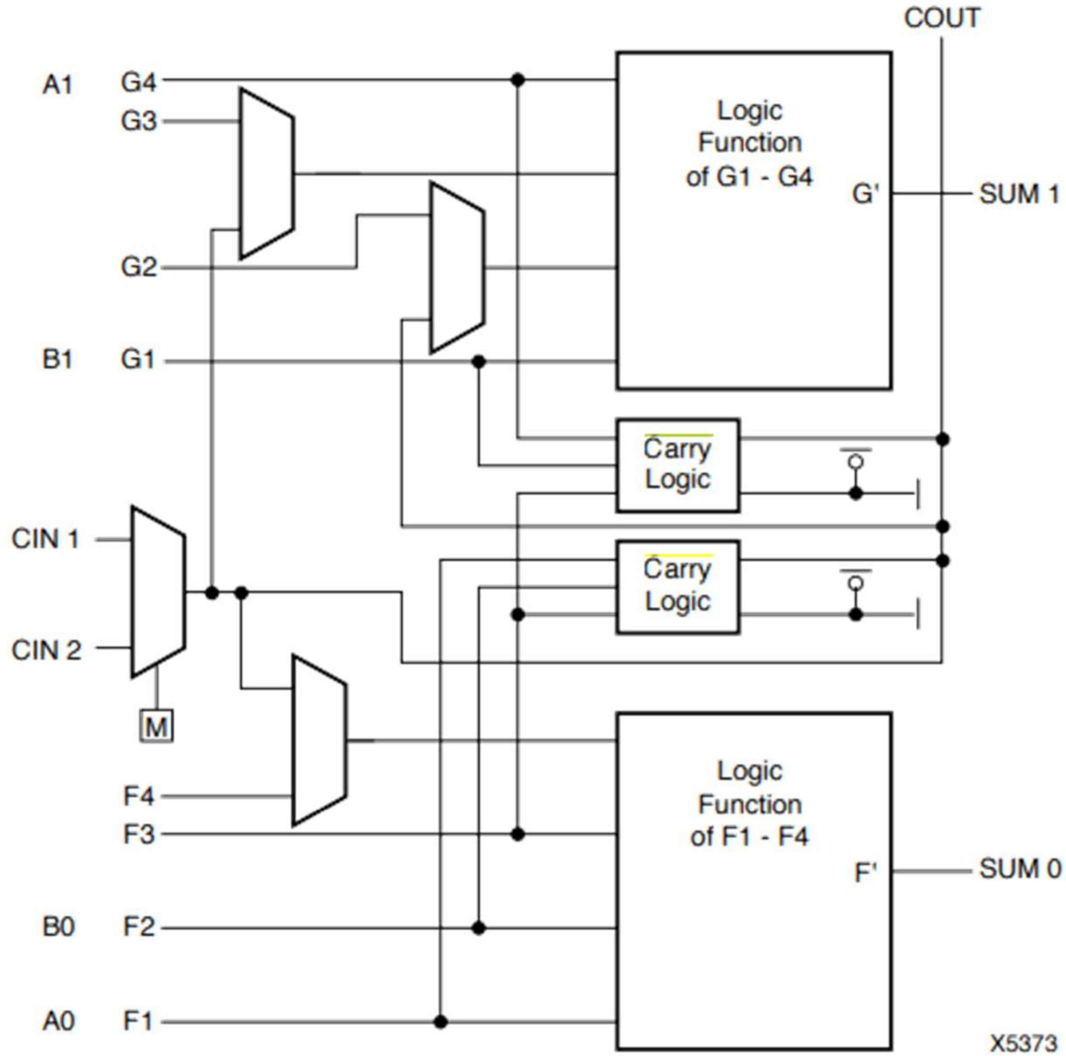

[XC4000, XC4000A, XC4000H Logic Cell Array Families]

A 16-bit adder requires nine CLBs and has a combinatorial carry delay of 20.5 ns. Compare that to the 30 CLBs and 50 ns, or 41 CLBs and 30 ns in the XC3000 family.

Hard Multipliers (2000s)

- Motivating forces
  - DSP became an important domain
  - very expensive and slow to multiply in LUTs
  - dies large enough to spare some area
- Virtex-II hardwired multiplier “macro” blocks
  - 18-bit inputs, full 36-bit product
  - explicit instantiation or inferable from RTL
  - relatively cheap (since native implementation)
  - Still no hard adders

Adders came later as a part of MAC in DSP slices

In the meanwhile, multiply faster/cheaper than add!!

An Early Multiplier Blocks Xilinx Virtex-II, circa 2000

[Virtex-II Pro and Virtex-II Pro X Platform FPGAs]

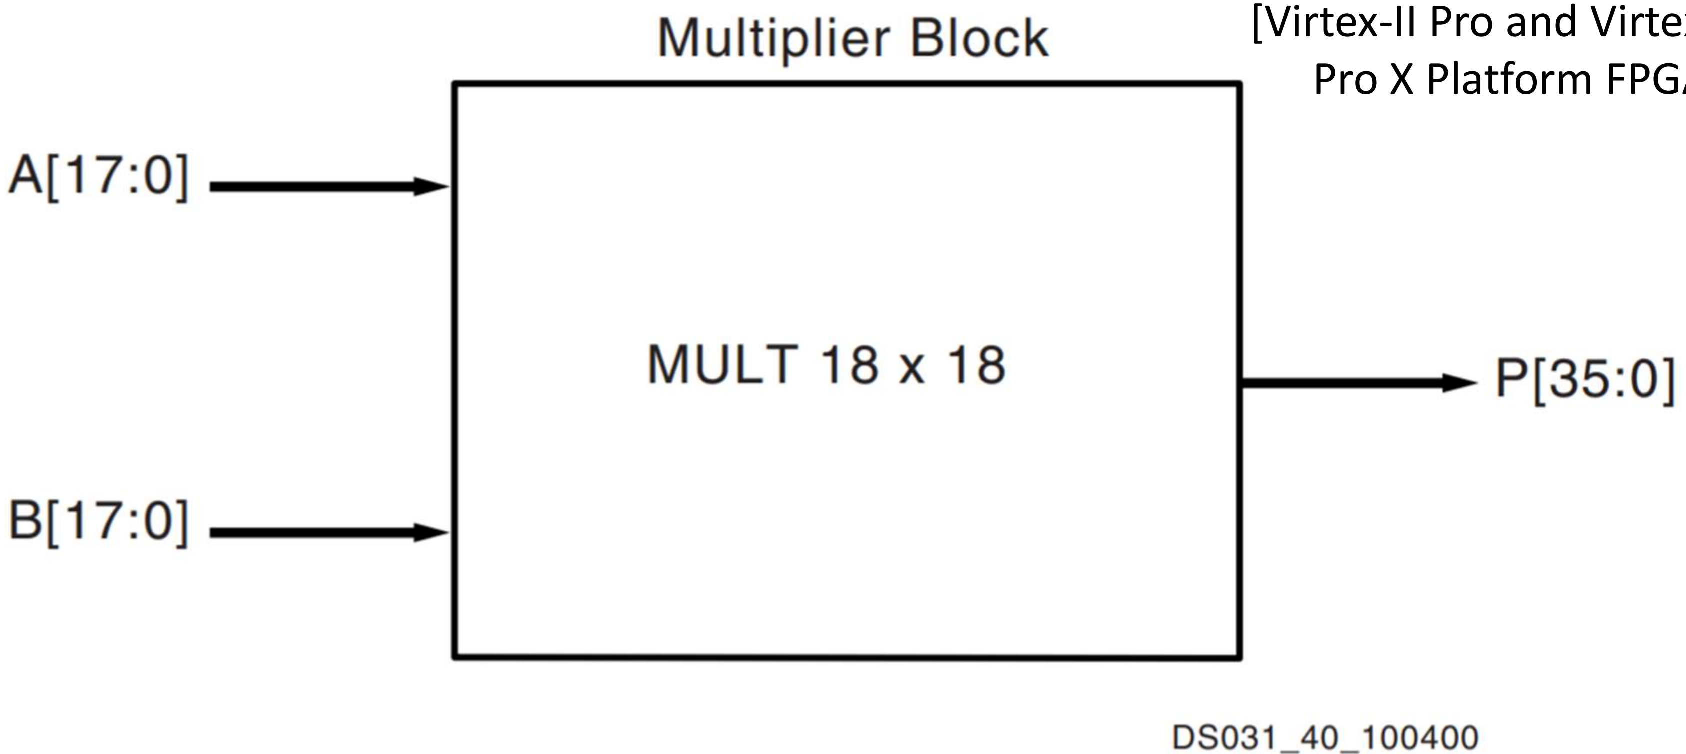

Figure 54: **Multiplier Block**

Where are these hard DSP slices?
How to get to them?

Ultrascale DSP48E2

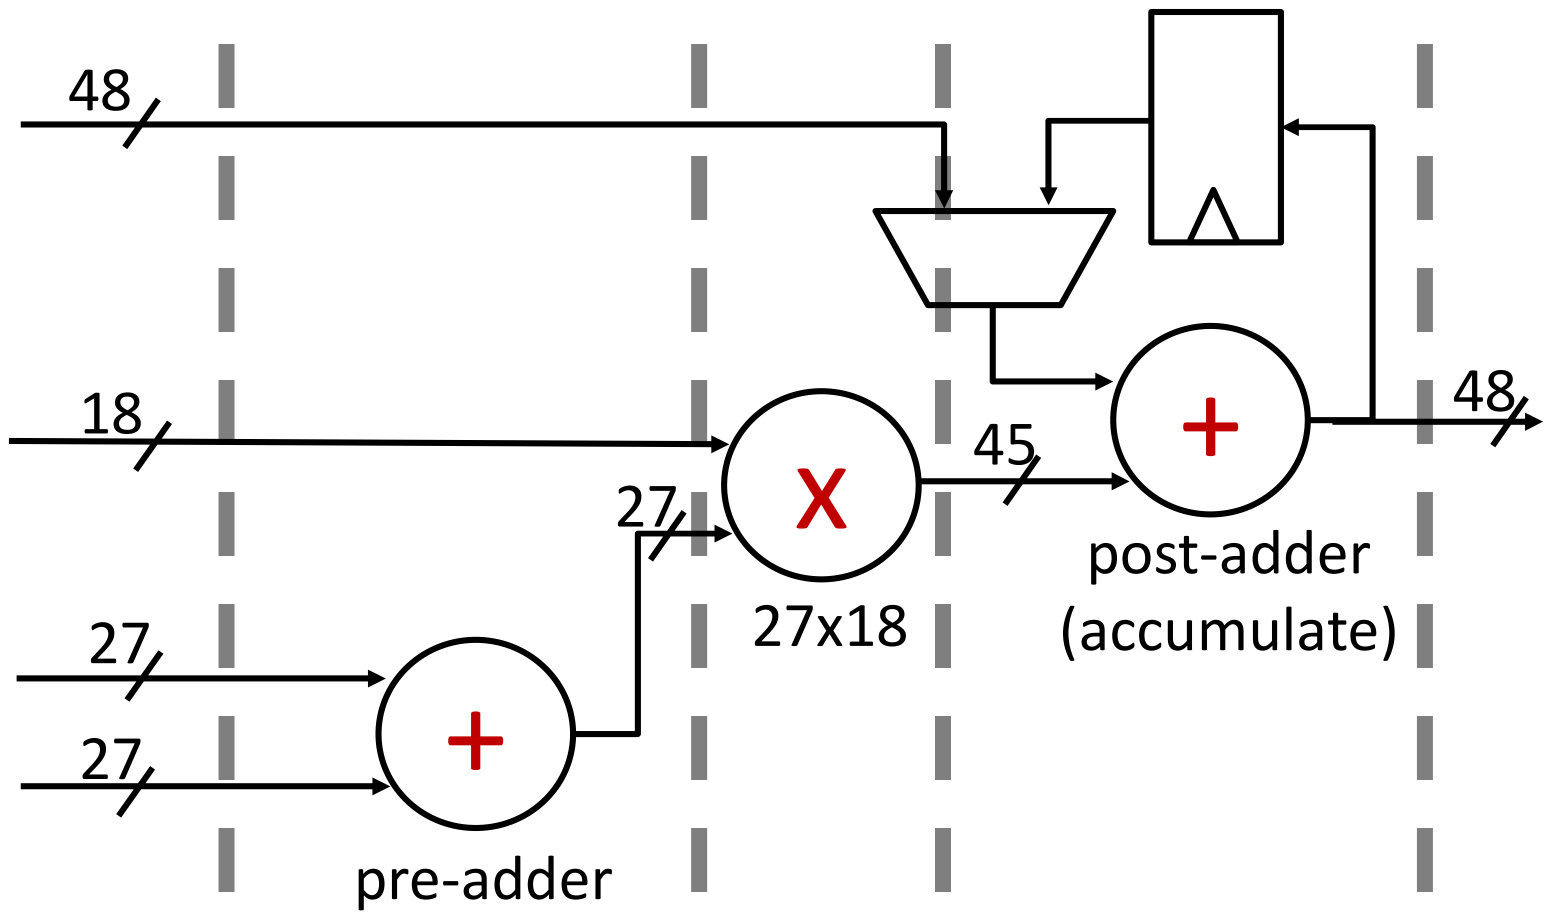

optional pipeline stages
inferable from RTL and retiming

Aside: Register Retiming

- Local transformations

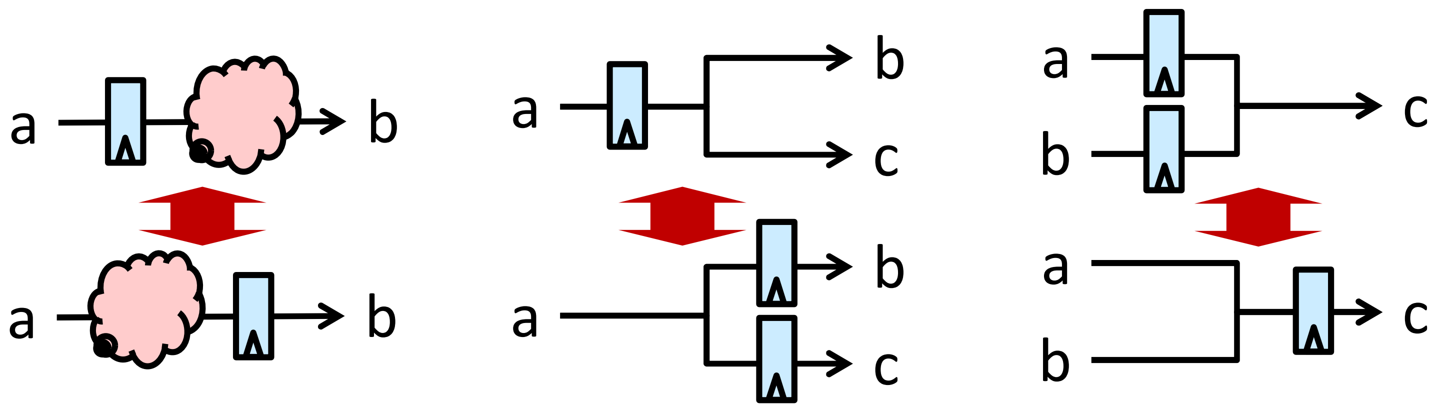

- Preserves I/O relationships
- Tools use retiming
  - balance critical paths
  - absorb FFs into hard macros

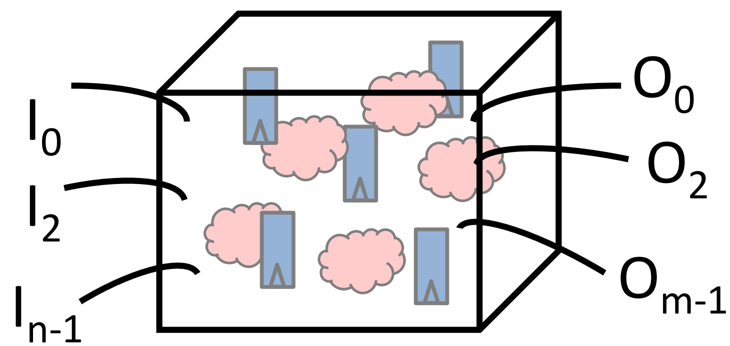

Pipelined multiply ➔

```
always@(posedge clk) begin
   a1<=a; b1<=b;
   a2<=a1; b2<=b1;
   c<=a2*b2;
end
```

Stratix/Arria-10 IEEE-754 DSPs

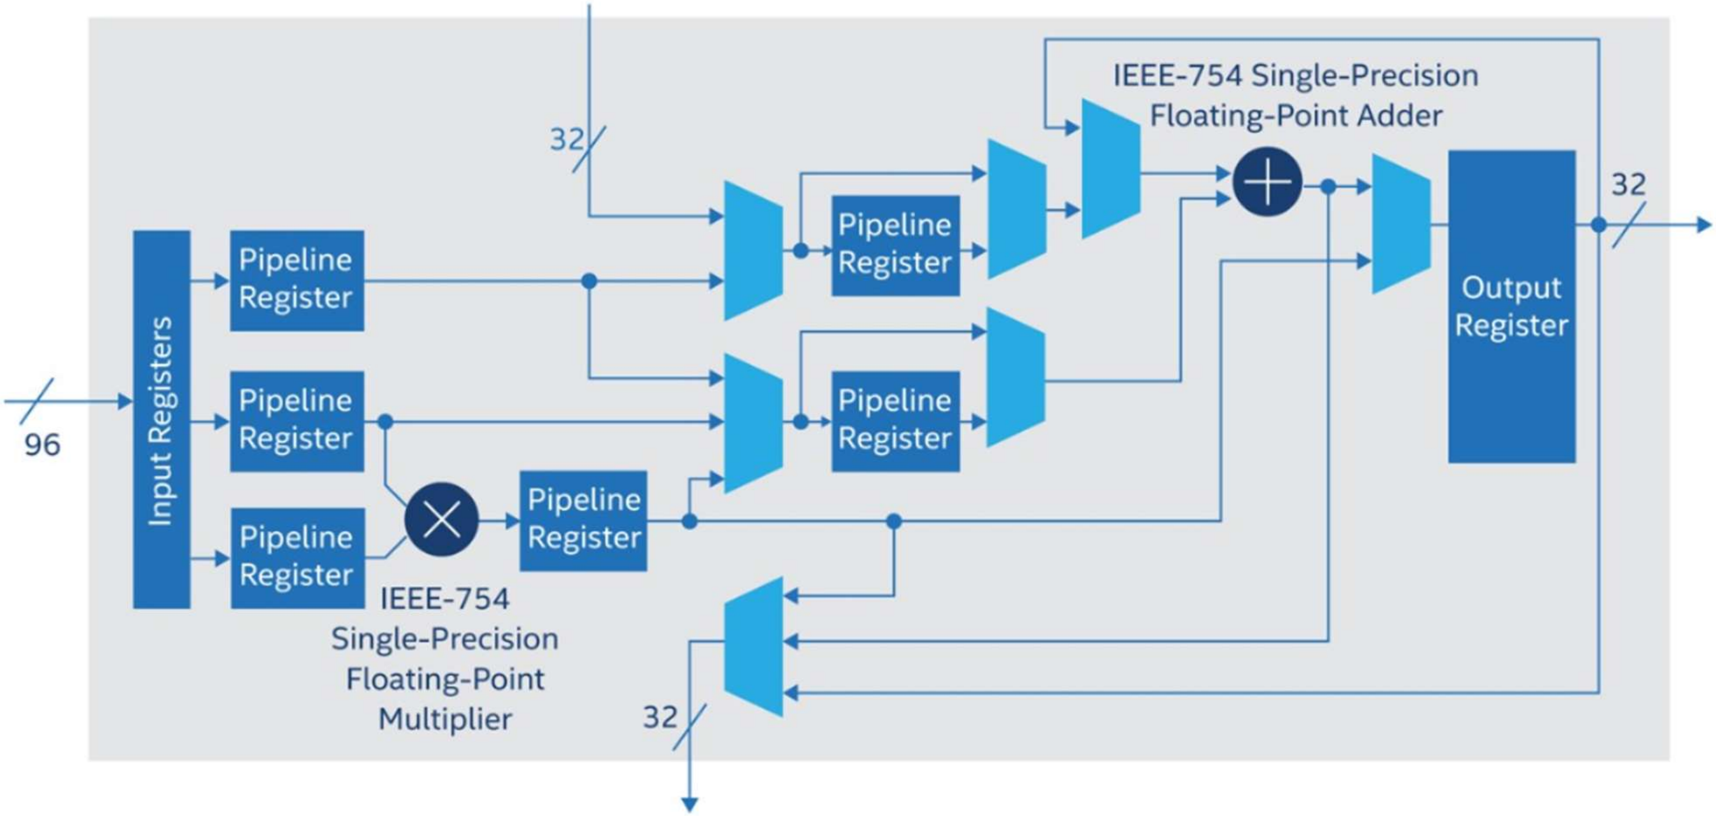

Intel® Stratix® 10 Device DSP Block: Single-Precision Floating Point

[Intel Stratix-10 FPGA Features]

Memory

- Flip-flops relatively scarce (only 1-bit per CLB)
- Need more storage when applications moved beyond FSM controllers and glue logic
- Option A: LUTs repurposable as 16x1-bit SRAMs
- Option B: 4Kb (now 32Kb) 2-ported SRAM blocks
  - very compact, very fast because native in silicon
  - explicit instantiation or inferable from RTL
    (tool can even decide which SRAM option to use)
  - configurable and combinable to a wide range of sizes and aspect ratios

Where are they? How to connect up to them?

MACROs: a disturbance in the force . . .

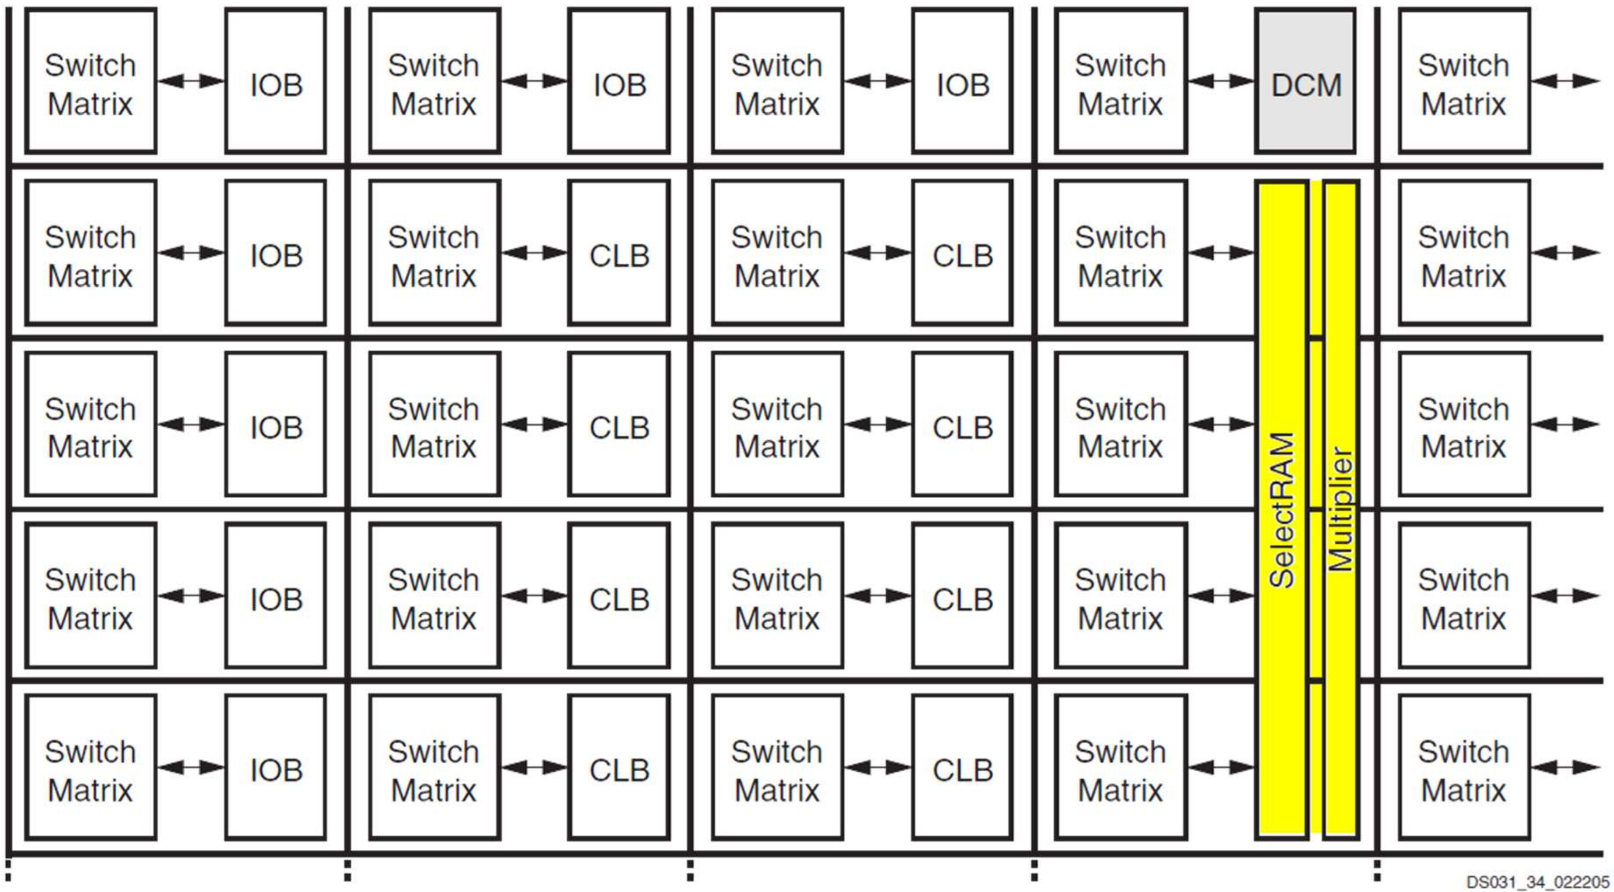

Too much vs not enough?

Benefit of using macro outweigh cost of getting to one?

[Figure 48: Virtex-II Platform FPGAs: Complete Data Sheet]

FPGA fabric not true blank slate

- FPGA Macros (especially RAM and DSP)
  - coarse functions and structures
  - some powerful but arbitrarily specific features
  - penalty is too huge to not get it right
- Inferable from RTL but . . . .
  - hard macros only does what it does
  - tools cannot recognize all "functional equivalent" descriptions
  - good idea to check inference report

Straight out-of-the-box ASIC RTL likely suboptimal, sometimes not-mappable

Example: Flip-Flops Inference

- Use asynch set or reset
  - ⇒ not all FFs have async reset; prevents DSP retiming
- Use both set and reset
  - ⇒ no FF has this; emulated externally with LUTs
- Use set and reset operationally
  - ⇒ set/reset cannot use special global lines
- Active-low set/reset and enables
  - ⇒ need LUTs to turn active-high

How could you know this?

- BRAM cannot be used if combinational read
- Shift registers can be made out of LUTs

  BUT! no set/reset and can't read middle bits
- Registers will retime into multiplier and DSP (if no asynch reset)
- Use "initial" for power-on reset
- Timing analysis doesn't do "latches"
- Many, many more like this. . .

Always want to RTMF!

Processor Cores

- Not everything needs to be in hardware; not everything improves when made into hardware
- Augment fabric with simple embedded CPUs
  - provide universality of functionality
  - easy handling of irregular, sequential operations
  - easy handling anything that doesn't need to be fast
- Interests developed in early 2000s when FPGA applications grew to systems with DRAM, video, and Ethernets, etc.

Hard or soft core?

Hardcore vs Softcore

- First came PowerPC hardcores on Virtex-II
  - you got 2 whether you needed it or not
  - new tool promote IP-based system building
  - entirely soft-logic built surroundings: busses and IPs (DRAM controller, Ethernet, video, ...)
- Microblaze softcores took over in later rounds
  - Xilinx proprietary ISA (runs OS, gcc and all that)
  - configurable for cost-performance tradeoff
  - available in RTL to some folks
  - by this time, softcore footprint and performance was acceptable

Several 3rd-party softcores existed in that era, e.g., LEON SPARC

Embedding PowerPC in Fabric

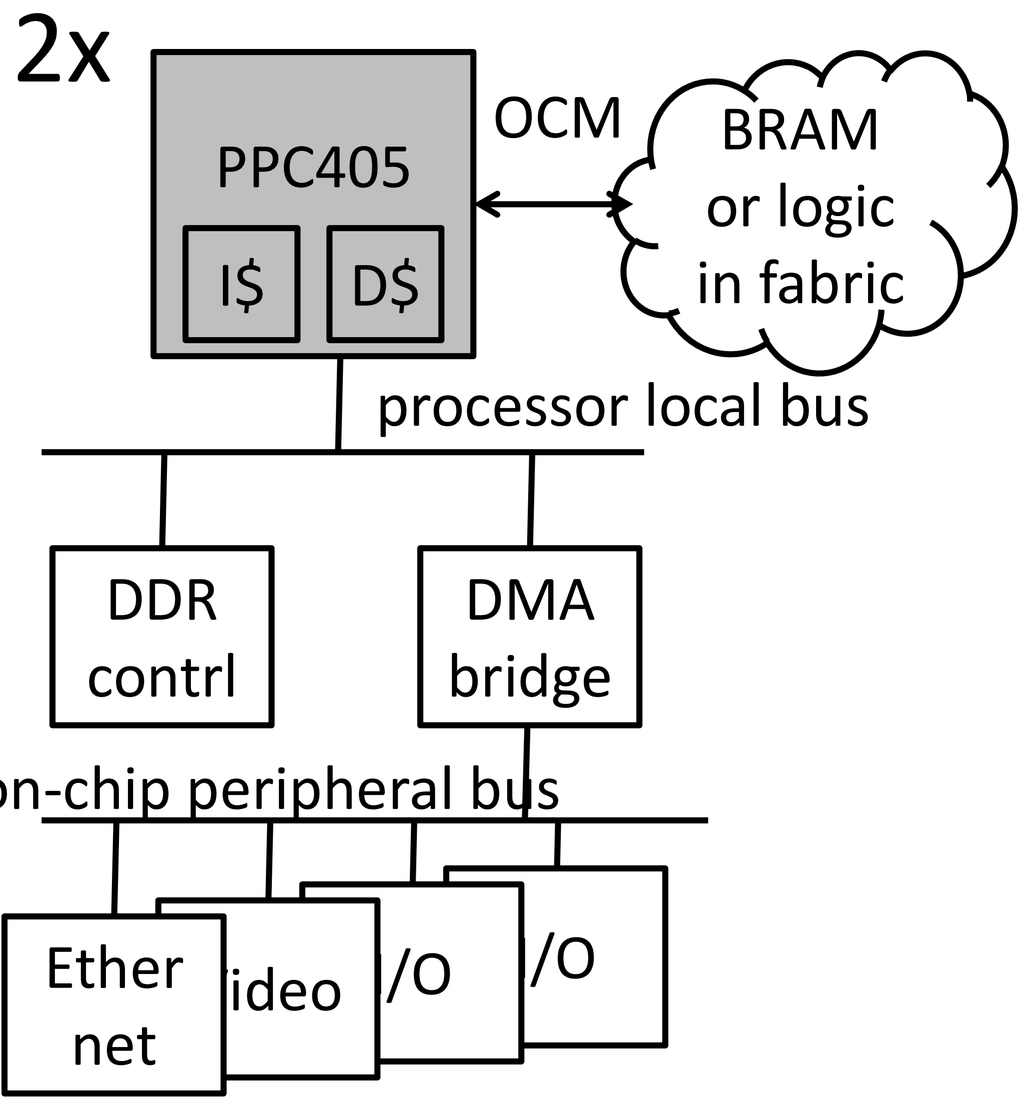

- everything else is soft
- two hierarchies of soft-logic busses

  (slow and slower)
- special on-chip memory (OCM) port allows ld/st directly into fabric
- CoreGen Library of IPs to hang off the busses

[Xilinx Vertex II, early 2000]

Hardcores Return in Virtex7 (~2010)

- This time in a complete, full-speed, fully-capable, two-core Cortex-A9 system
- Latest Ultrascale uses 64-bit ARMv8 Cortex-A53 + ARM R5 + Mali GPU
- Why ARMs?

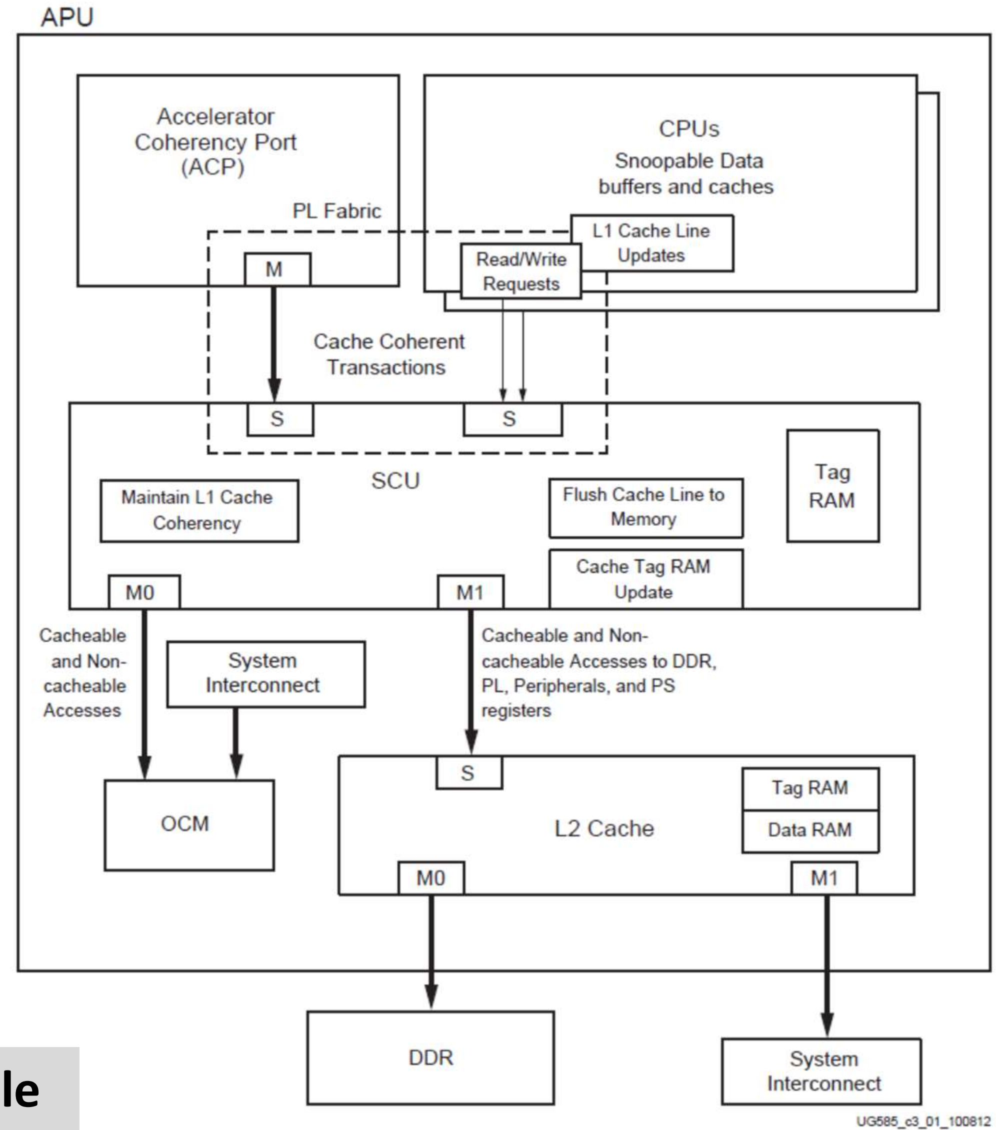

Figure 3-1: **APU Block Diagram**

[Figure 3-1, Zynq-7000 All Programmable SoC Technical Reference Manual]

Hardcore vs Softcore

- Table 4.2: The Zynq Book

Processor	Configuration	DMIPs
MicroBlaze 900LUT/700FF/ 2BRAM to 3800LUT/3200FF/ 6DSP/21BRAM	area optimized (3-stage)	196
	perf. optimized (5-stage) with branch optimizations	228
	perf. optimized (5-stage) without branch optimizations	259 ??from book
ARM Cortex-A9	1GHz; both cores combined	5000

- Table 4.3: The Zynq Book

Processor	Configuration	CoreMark
MicroBlaze	125MHz; 5-stage (Virtex-5)	238
ARM Cortex-A9	1GHz; both cores combined	5927
ARM Cortex-A9	800MHz; both cores combined	4737

PPC405 about 1/5 of ARM in Figure 4.3 of The Zynq Book

Die Area "Return on Investment"

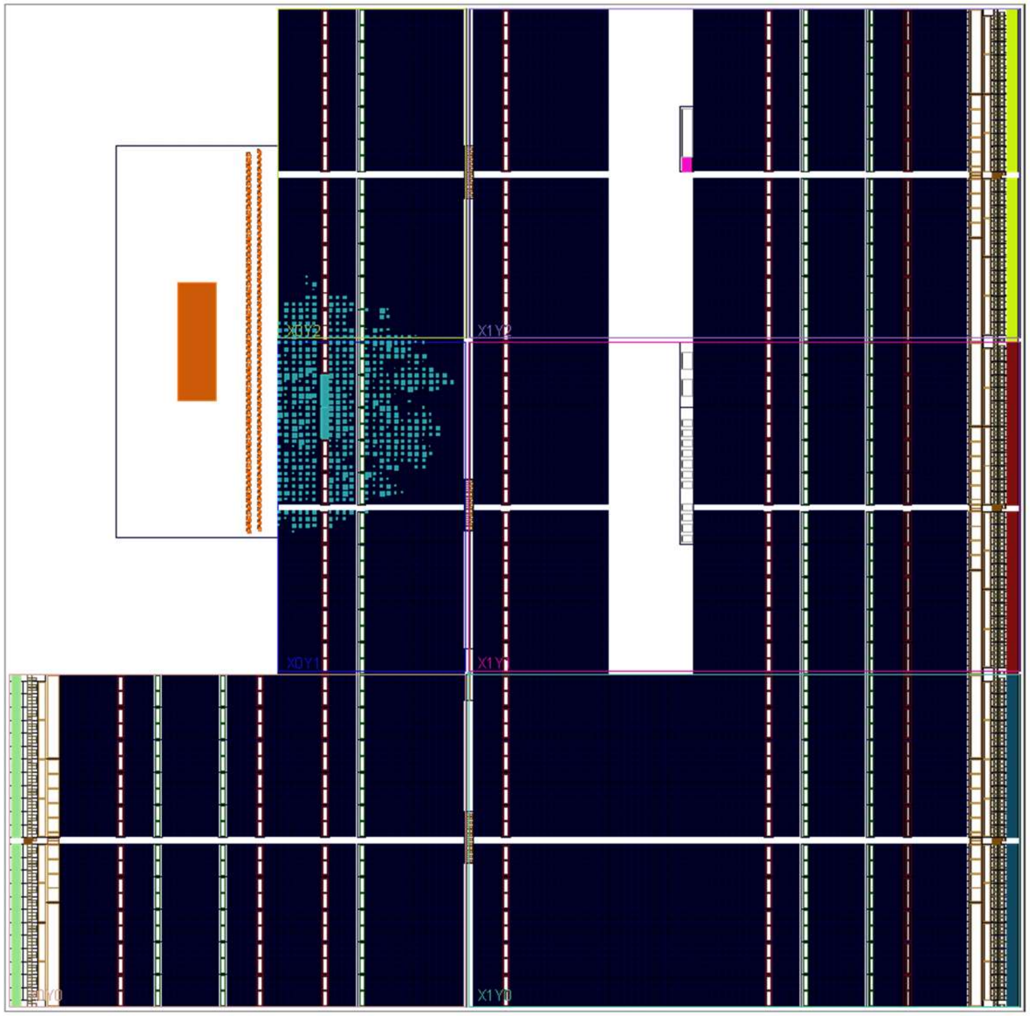

[Vivado screenshot XC7z020]

Soft-logic logic dominates die area, but compute/storage concentrated in DSP and BRAM—consider what if 100% soft or 100% hard

Xilinx ASMBL Architecture
(Application Specific Modular Block Arch.)

- Xilinx fabric assembled from composable tall-and-thin strip types, CLB, BRAM, DSP, I/O, etc.
- Derivative products at the cost of just new masks
  - vary capacity by composing more or less strips
  - domain-specialization by varying ratios of strips e.g., {DSP+IP} vs logic for DSP vs ASIC replacement market
  - variations handled by parameterization in design tool algorithms

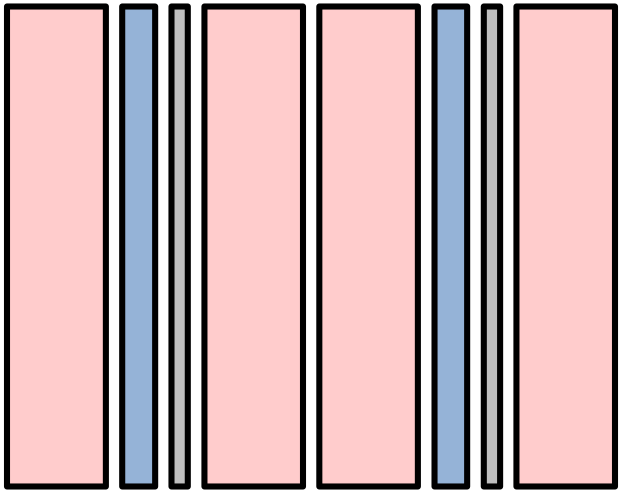

Stacked Silicon Interconnect (SSI)

- 2.5D stacking: multiple dies on passive interposer
  - lower latency, higher bandwidth, lower power than crossing package
  - much better yield than equivalent capacity monolithic device
  - mix dies for domain-specialization
  - possible to insert customer proprietary dies?

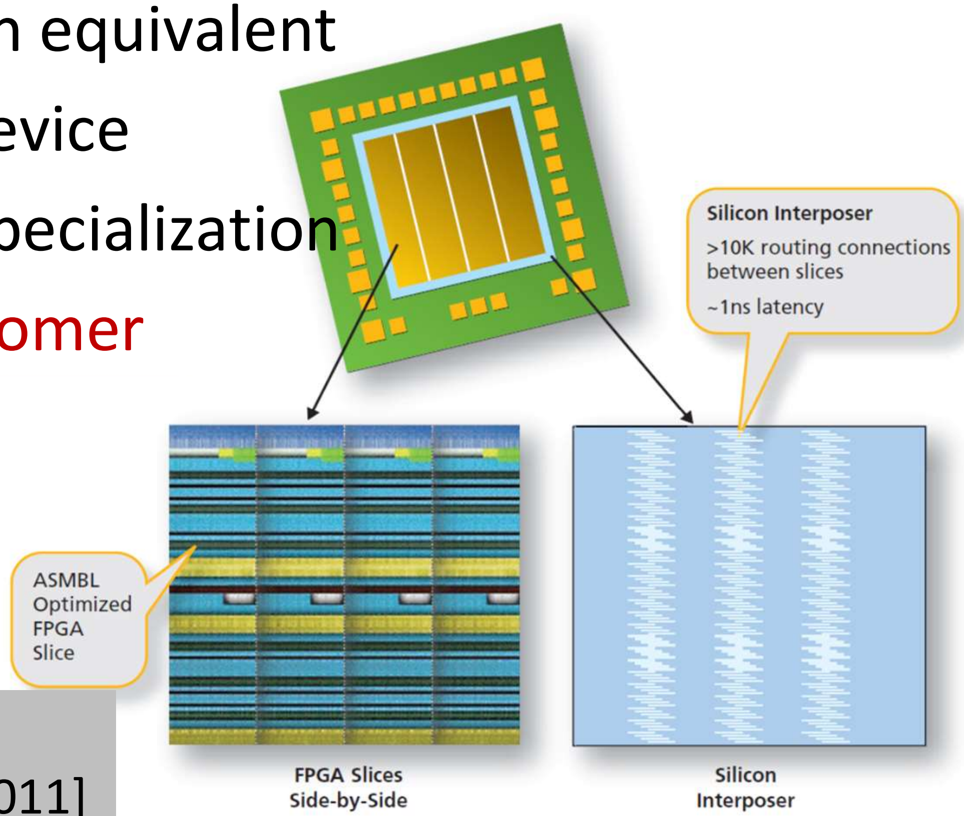

[Figure 1, Stacked & Loaded: Xilinx SSI, 28-Gbps I/O Yield Amazing FPGAs, Xcell, Q1 2011]

Intel's take on 2.5D with EMIB

Figure 8. Enhanced Flexibility and Scalability with Separate Transceiver Tiles

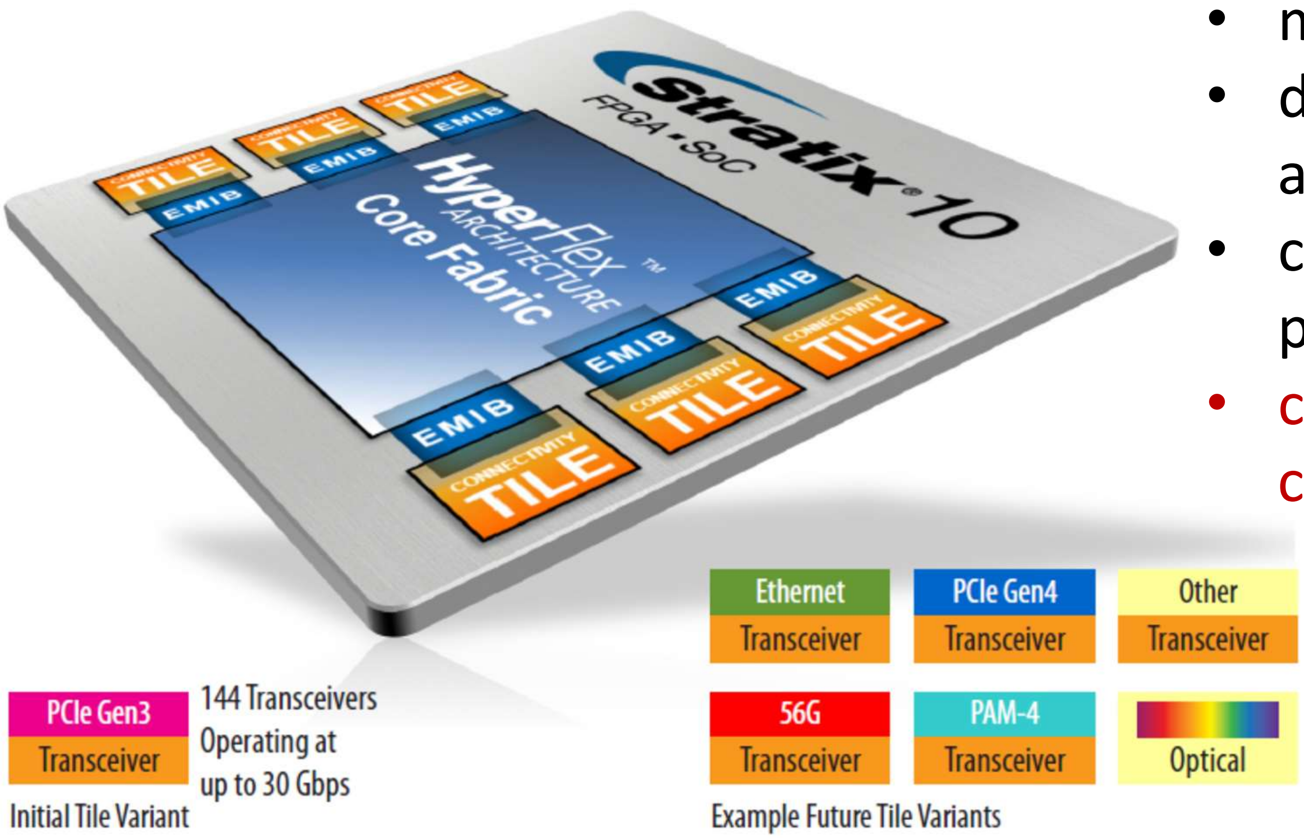

- monolithic fabric
- displace noisy, hot analog IPs
- connect same-package HBMs
- connect 3rd-party chiplets?

[Figure 8, Enabling Next-Generation Platforms Using Altera's 3D System-in-Package Technology]

Reviewing Hard IPs Added Over Time

- 1990s
  - fast carry
  - LUT RAM
  - block RAM
- 2000s
  - programmable clock generator
  - PowerPC core
  - gigabit transceiver
  - multiplier and DSP splices
  - Ethernet and PCI-E
- 2010s
  - system monitor
  - ADC
  - power management
  - ARM cores and GPU
  - DRAM controller
  - floating point arithmetic
  - **"UltraRAM" hierarchy (up to 500Mbits)**
  - **HBM controllers**
- 2020s . . . . *next lecture*

Chicken or Egg First?

- 1990s: glue logic, embedded cntrl, interface logic
  - reduce chip-count, increase reliability
  - rapid roll-out of “new” products
- 2000s: DSP and HPC
  - strong need for performance
  - abundant parallelism and regularity
  - low-volume, high-valued
- 2010s: communications and networking
  - throughput performance
  - fast-changing designs and standards
  - price insensitive
  - $value in field updates and upgrades

SoC with reconfigurable fabric (2010s)

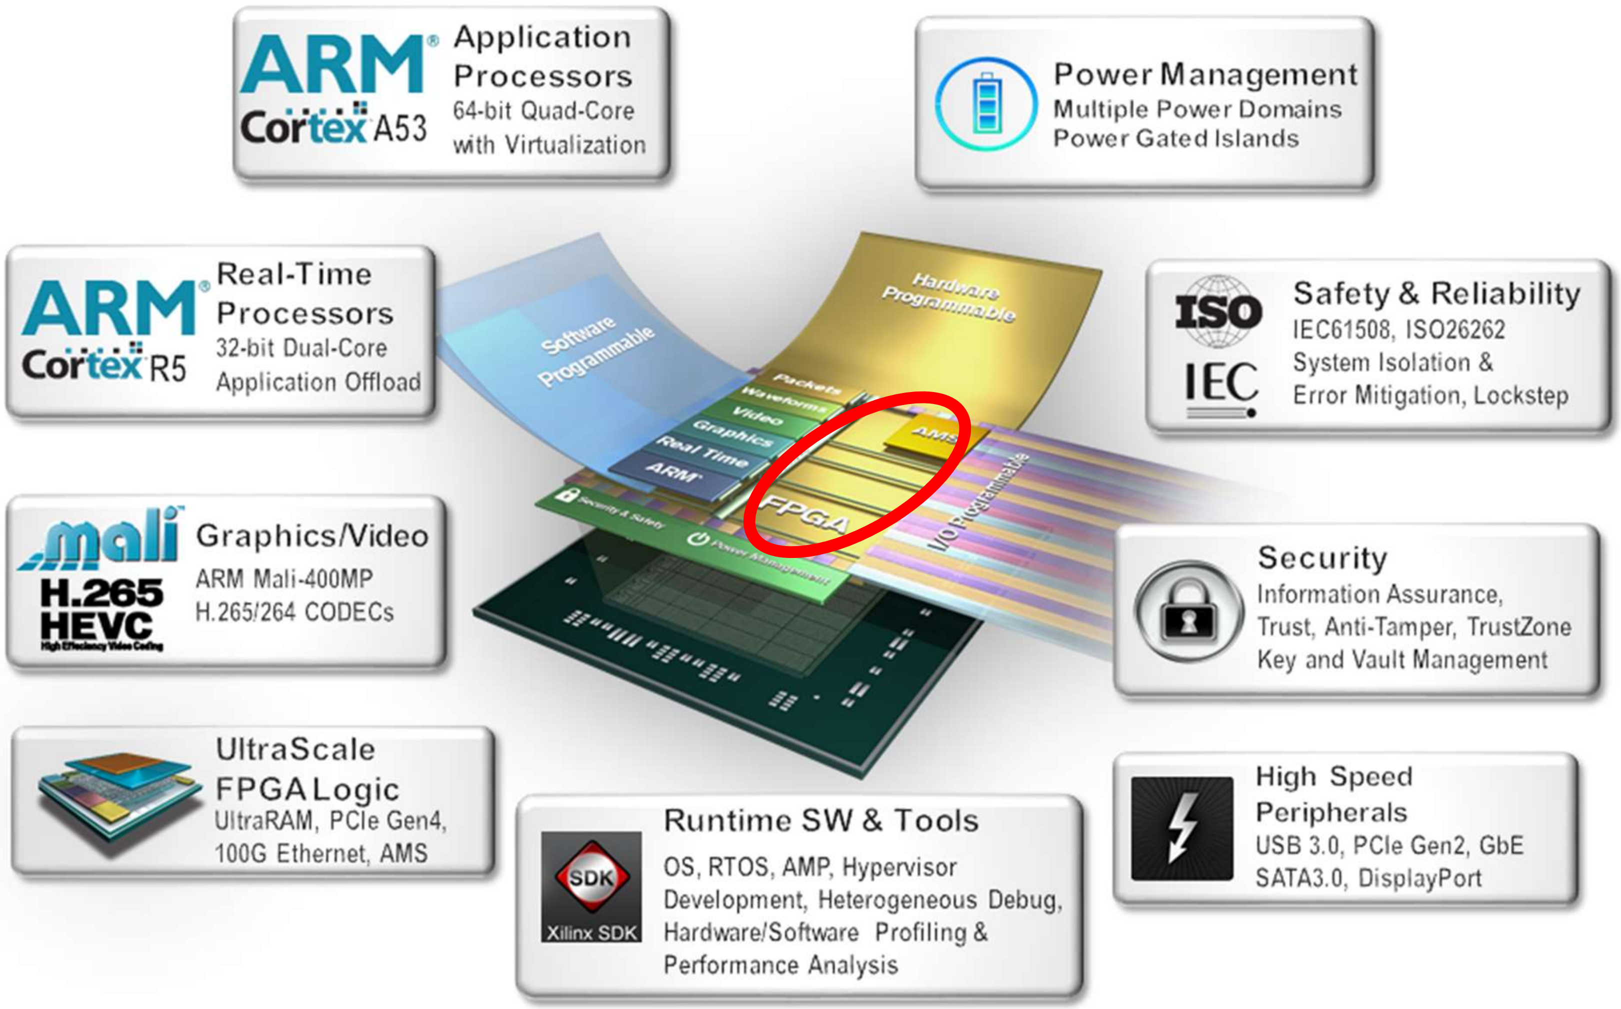

[http://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale-mpsoc.html]

Xilinx Vertex Ultrascale Offerings

	Device Name	XCVU065	XCVU080	XCVU095	XCVU125	XCVU160	XCVU190	XCVU440
Logic Resources	System Logic Cells (K)	783	975	1,176	1,567	2,027	2,350	5,541
	CLB Flip-Flops	716,160	891,424	1,075,200	1,432,320	1,852,800	2,148,480	5,065,920
	CLB LUTs	358,080	445,712	537,600	716,160	926,400	1,074,240	2,532,960
Memory Resources	Maximum Distributed RAM (Kb)	4,830	3,980	4,800	9,660	12,690	14,490	28,710
	Block RAM/FIFO w/ECC (36Kb each)	1,260	1,421	1,728	2,520	3,276	3,780	2,520
	Block RAM/FIFO (18Kb each)	2,520	2,842	3,456	5,040	6,552	7,560	5,040
	Total Block RAM (Mb)	44.3	50.0	60.8	88.6	115.2	132.9	88.6
Clock Resources	CMT (1 MMCM, 2 PLLs)	10	16	16	20	28	30	30
	I/O DLL	40	64	64	80	120	120	120
	Transceiver Fractional PLL	5	8	8	10	13	15	0
I/O Resources	Maximum Single-Ended HP I/Os	468	780	780	780	650	650	1,404
	Maximum Differential HP I/O Pairs	216	360	360	360	300	300	648
	Maximum Single-Ended HR I/Os	52	52	52	52	52	52	52
	Maximum Differential HR I/O Pairs	24	24	24	24	24	24	24
Integrated IP Resources	DSP Slices	600	672	768	1,200	1,560	1,800	2,880
	System Monitor	1	1	1	2	3	3	3
	PCIe® Gen1/2/3	2	4	4	4	4	6	6
	Interlaken	3	6	6	6	8	9	0
	100G Ethernet	3	4	4	6	9	9	3
	GTH 16.3Gb/s Transceivers	20	32	32	40	52	60	48
	GTY 30.5Gb/s Transceivers	20	32	32	40	52	60	0
Speed Grades	Commercial	–	–	–	–	–	–	-1
	Extended	-1H -2 -3	-1H -2 -3	-1H -2 -3	-1H -2 -3	-1H -2 -3	-1H -2 -3	-2 -3
	Industrial	-1 -2	-1 -2	-1 -2	-1 -2	-1 -2	-1 -2	-1 -2

[UltraScale FPGA Product Tables and Product Selection Guide (XMP102)]

Intel Agilex-10 Offerings

	PRODUCT LINE	AGF 006	AGF 008	AGF 012	AGF 014	AGF 019	AGF 022	AGF 023	AGF 027
Resources	Logic elements (LEs)	573,480	764,640	1,178,525	1,437,240	1,918,975	2,208,075	2,308,080	2,692,760
	Adaptive logic modules (ALMs)	194,400	259,200	399,500	487,200	650,500	748,500	782,400	912,800
	ALM registers	777,600	1,036,800	1,598,000	1,948,800	2,602,000	2,994,000	3,129,600	3,651,200
	High-performance crypto blocks	0	0	0	0	2	0	2	0
	eSRAM memory blocks	0	0	2	2	1	0	1	0
	eSRAM memory size (Mb)	0	0	36	36	18	0	18	0
	M20K memory blocks	2,844	3,792	5,900	7,110	8,500	10,900	10,464	13,272
	M20K memory size (Mb)	56	74	115	139	166	212	204	259
	MLAB memory count	9,720	12,960	19,975	24,360	32,525	37,425	39,120	45,640
	MLAB memory size (Mb)	6	8	12	15	20	23	24	28
	I/O PLL	12	12	16	16	10	16	10	16
	Variable-precision digital signal processing (DSP) blocks	1,640	2,296	3,743	4,510	1,354	6,250	1,640	8,528
	18 x 19 multipliers	3,280	4,592	7,486	9,020	2,708	12,500	3,280	17,056
	Single-precision or half-precision tera floating point operations per second (TFLOPS)	2.5 / 5.0	3.5 / 6.9	6.0 / 12.0	6.8 / 13.6	2.0 / 4.0	9.4 / 18.8	2.5 / 5.0	12.8 / 25.6
	Maximum EMIF x72	2	2	4	4	2	4	2	4

84	240	384	240	384
2	4	4	4	4

DDR4, QDR IV

authentication, physically unclonable function (PUF),
ı, side-channel attack protection

.50 GHz with 32 KB I/D cache, NEON coprocessor, 1 MB L2
m memory management unit, cache coherency unit, hard
x3, UART x2, serial peripheral interface (SPI) x4, I2C x5,
er x4

or bifurcateable 2x PCIe 4.0 x8 (EP) or 4x 4.0 x4 (RP)
t 32 Gbps (NRZ) /12 channels at 58 Gbps (PAM4) - RS &

/25/50/100/200/400 GbE FEC/PCS/MAC)
25/50/100/200 Gbps FEC/PCS)
IEEE 1588 v2 support
PMA direct

Tile Resources	E-Tile	Transceiver channel count : Up to 24 channels at 28.9 Gbps (NRZ) / 12 channels at 58 Gbps (PAM4) - RS & KP FEC[1] Networking support : - 400GbE (4 x 100GbE hard IP blocks (10/25 GbE FEC/PCS/MAC)) IEEE 1588 v2 support PMA direct
	P-Tile	PCIe hard IP block (4.0 x16) or bifurcateable 2x PCIe 4.0 x8 (EP) or 4x 4.0 x4 (RP) SR-IOV 8PF / 2kVF VirtIO support Scalable IOV

[Intel FPGA Product Catalog]

Today's Diverging Architectures

Are they FPGAs?

- spatial data/compute
- highly concurrent
- finely controllable
- reprogrammable

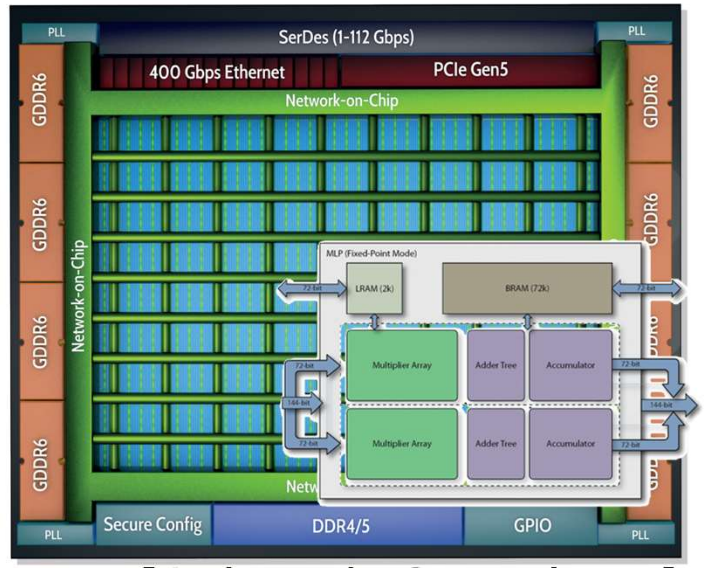

[Achronix Speedster]

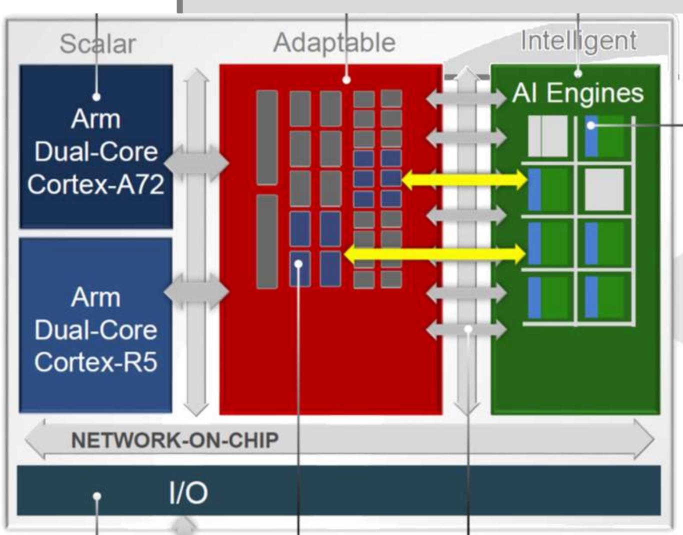

[Xilinx Versal]

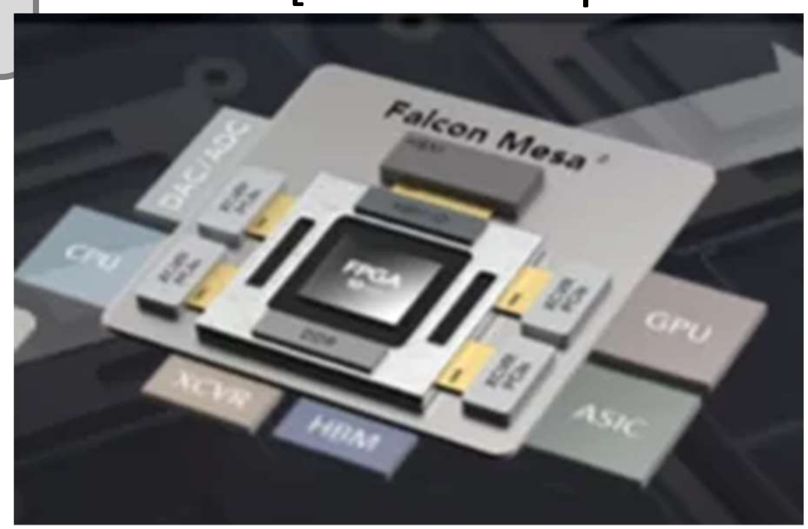

[Intel Agilex]

Parting Thoughts

- FPGAs steadily moved away from universal fabric
  - efficiency of hardwired logic (driven by application demands) complements flexibility of reconfig. logic
  - architected deliberately to play up this advantage
- Retain a high degree of regularity to ease design and manufacturing
  - fastest way to use up transistors from Moore's Law
  - power and performance advantage by just being first on new process
- Architectural evolution both push-and-pull with applications