

The Design of the KiloCore Chip

Aaron Stillmaker*, Brent Bohnenstiehl,
Bevan Baas

DAC 2017: Design Challenges of New Processor Architectures




University of California, Davis
VLSI Computation Laboratory
June 20, 2017



*Currently with California State University, Fresno

Outline

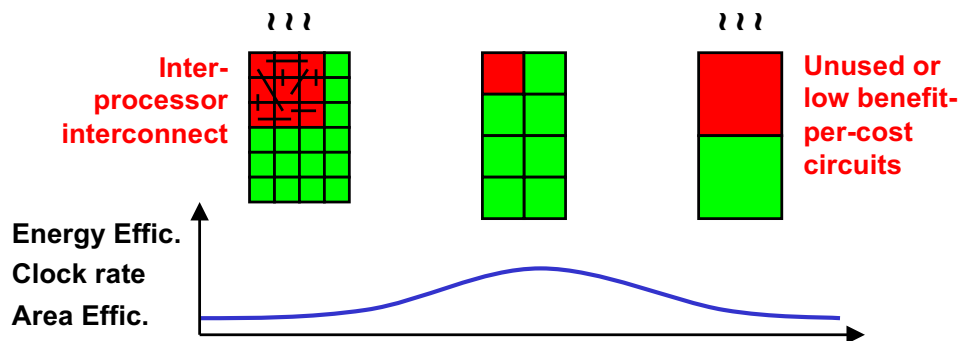
- Introduction and Motivation
 - Core Scaling Trend
- KiloCore Architectural Design
- Physical Chip Design
 - Chip Flow
 - Design Challenges
- Final Die
- Application Measurements



3

Optimal Computational Tile Size

- The most efficient implementations (energy, throughput, circuit area) have:
 - Processor sizes that capture computational kernels with few excess circuits



4

Key Properties of KiloCore

- **Very small processor tiles**
 - Vast numbers of processors per chip
 - Processors can be used for non-traditional purposes in highly efficient implementations
- Processors dissipate **very little energy per workload when active**
- Processors dissipate **very low power when idle**
- Essentially **no algorithm-specific hardware** in the programmable processors

5

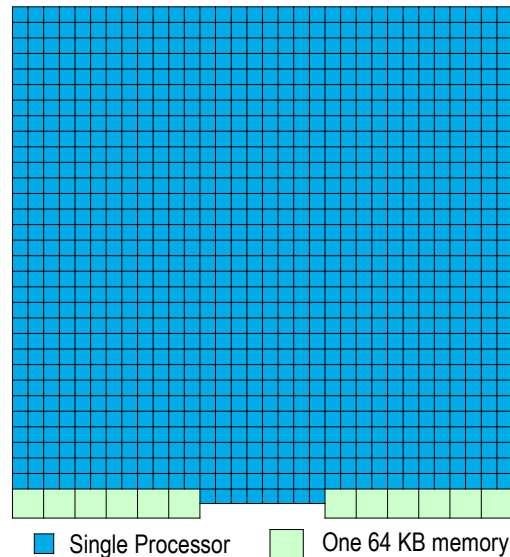
Outline

- Introduction and Motivation
 - Core Scaling Trend
- **KiloCore Architectural Design**
- Physical Chip Design
 - Chip Flow
 - Design Challenges
- Final Die
- Application Measurements

6

KiloCore Design


- Contains exactly 1,000 processors on one chip
- One of the first fabricated chips to contain 1,000 processors
- Fastest clock rate processor designed at a university
- Didn't receive all libraries until **34 days before taping out**
- 12 memories containing 64 KB each for 768 KB of shared memory
 - Memories are accessible by two processors directly above each



7

GALS Clocking

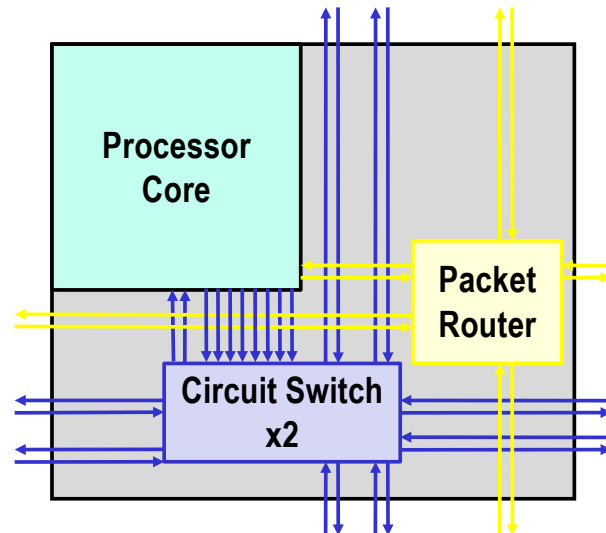
- KiloCore contains a fully-independent Globally-Asynchronous Locally-Synchronous (GALS) clock domain in each of its 1000 processors, 1000 packet routers, and 12 independent memories
 - Processor programmable clock oscillators are ~1% of tile area
 - Router oscillators are simplified and very small
- Each of the 2012 clock oscillators are placed inside their own clock domains—there are no global clock signals (except three for configuration and testing)
- Each clock oscillator is fully-unconstrained—oscillators may change frequency (below their f_{max}), halt, and restart arbitrarily to minimize power consumption
- Data transfer across clock domains is handled by dual-clock FIFOs



8

Inter-Processor Communication

- Source synchronous 16-bit communication
- Two layer source-synchronous circuit switched network
 - Up to 28 Gbps per link, 456 Gbps total tile I/O
- Single layer dynamic packet routing network
 - 9.1 Gbps maximum per port
 - 45.5 Gbps maximum



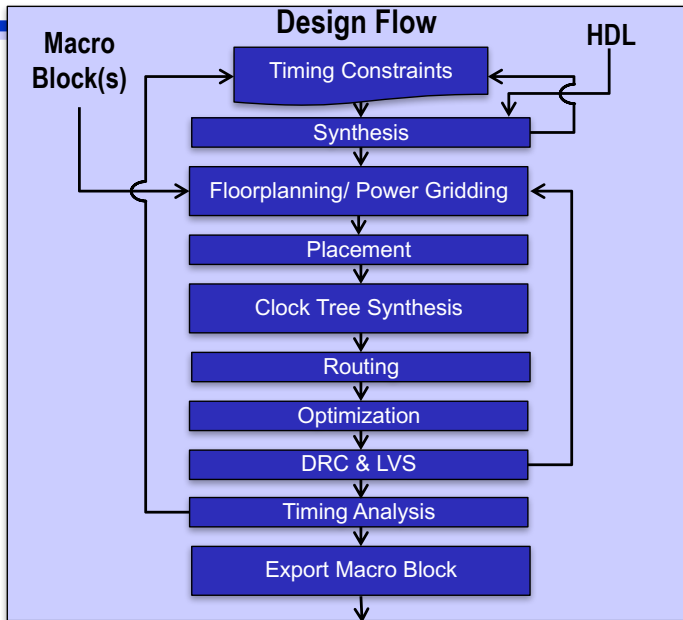
9

Outline

- Introduction and Motivation
 - Core Scaling Trend
- KiloCore Architectural Design
- **Physical Chip Design**
 - Chip Flow
 - Design Challenges
- Final Die
- Application Measurements

10

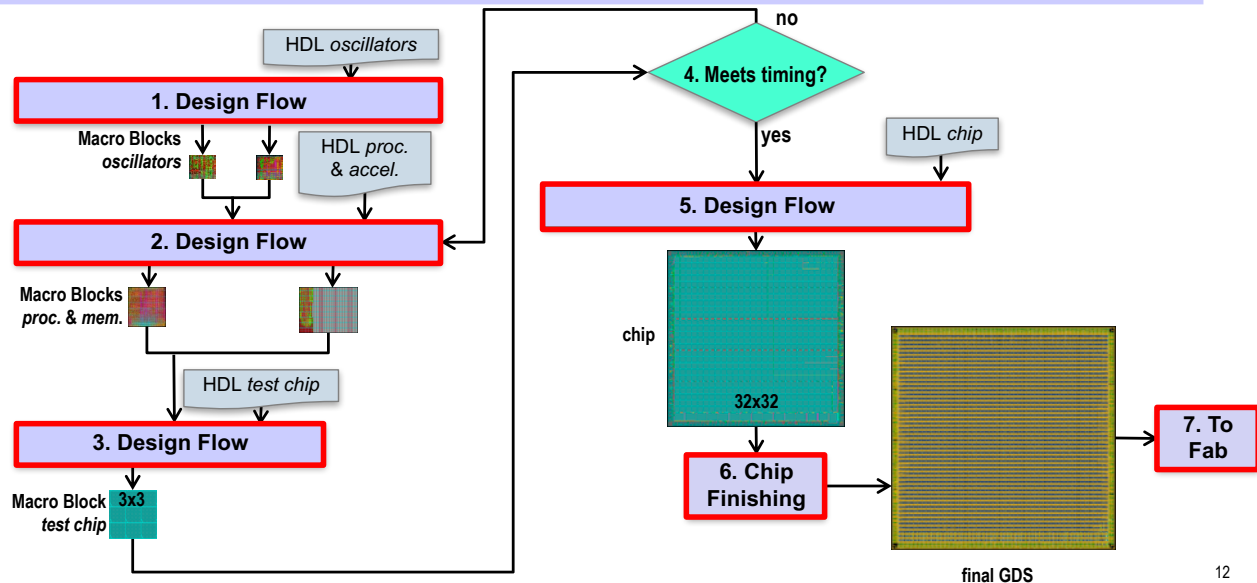
Bottom-Up Physical Design Flow



- Bottom-Up design flow
- Design aspects undecided when the physical design was started:
 - Routers
 - Inter-processor I/O
 - Chip I/O
 - Memories (both local and global)
 - Oscillators
 - Configuration
 - Number of processors

11

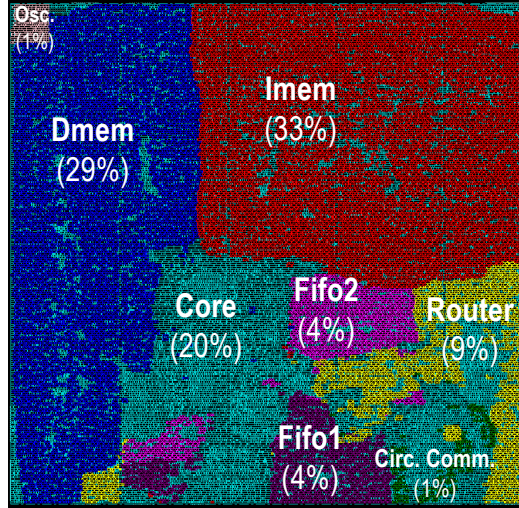
Bottom-Up Physical Design Flow



12

Processor Tile Implementation

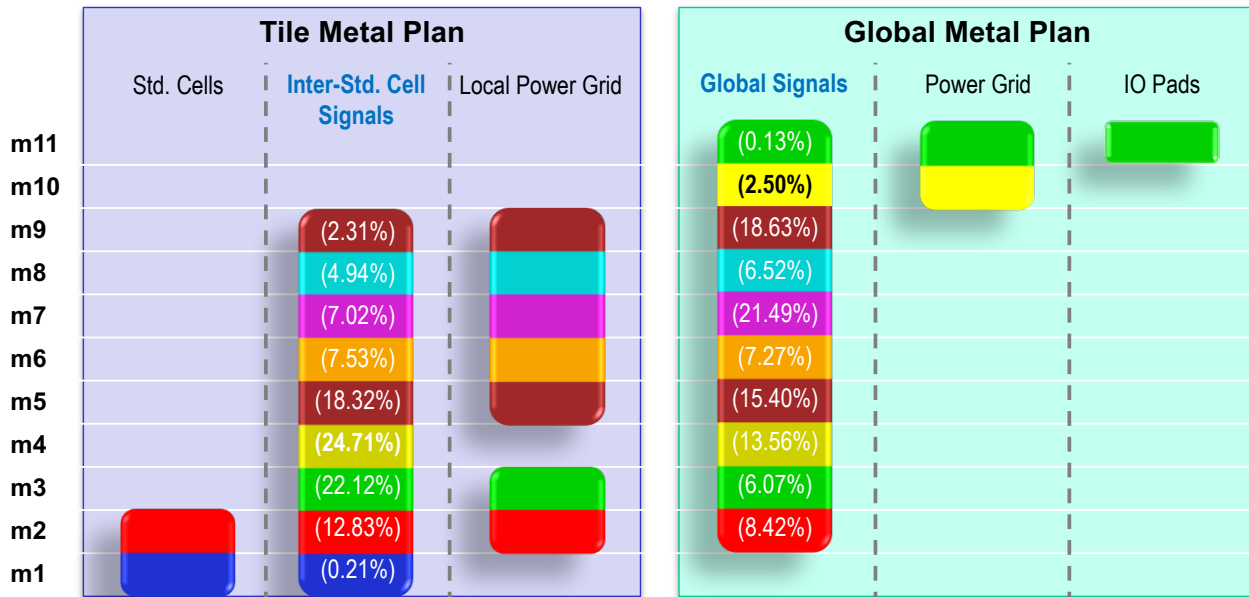
- Relatively simple processor design
 - No hard macros
 - Osc was pre-placed and routed
 - I/O pin locations specified for optimal abutment
 - Power rails specified
- Quick design iterations
 - Easily changed design aspects such as memory size or tile size
 - Verilog was changing up until 5 days before tapeout
- ~580k transistors per processor tile
- 239 μm “wide” × 232.3 μm “tall”



Single processor with highlighted gates (percentage of tile area)

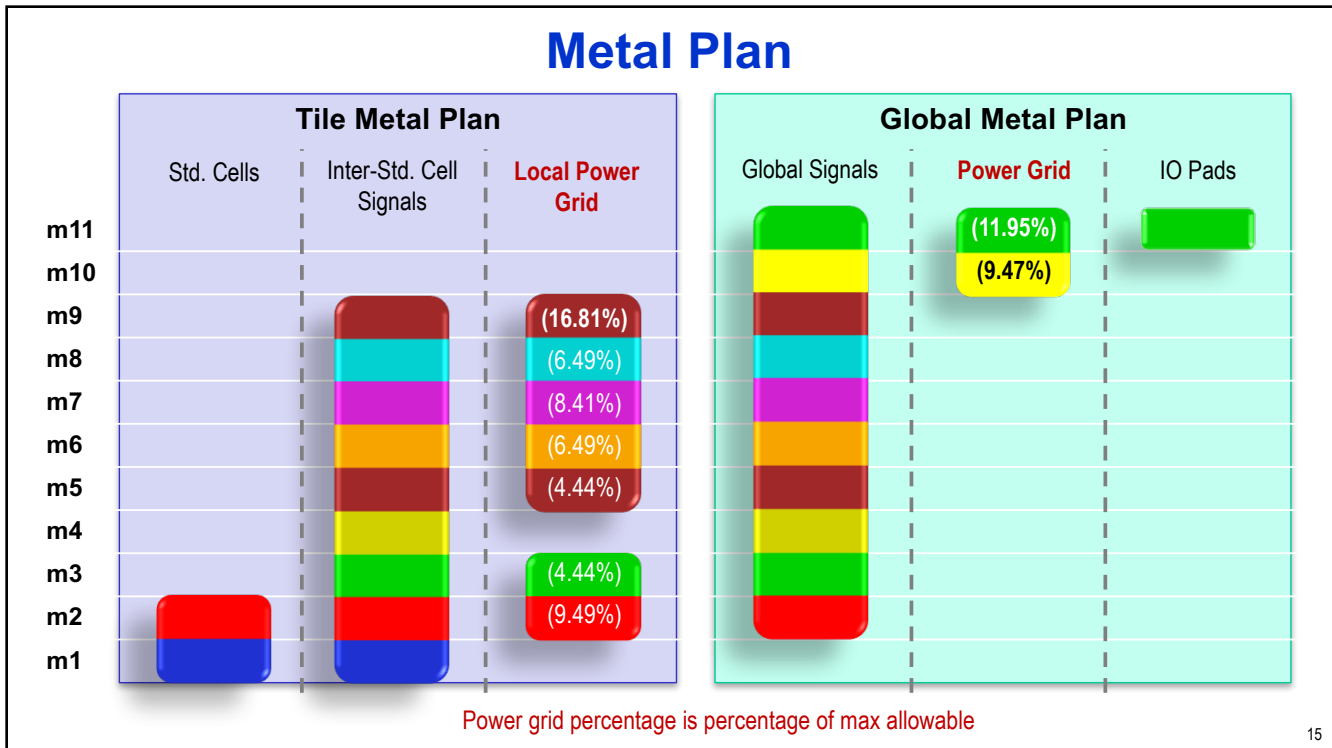
13

Metal Plan




Signal percentages are length per layer, per total signal length

14



15

Processor Level Power Rails




View focused on single core, showing array painted power

- Used 7 metal layers for local power grid
 - 39.6 μm horizontal pitch
 - 20.0 μm vertical pitch
- Connected local grid between all processors
 - Including the standard cell level power rails
- Manually drew in the grid of metal at array level

16

Global Power Rails

- Metal layers 10 and 11 are used for global power grid
 - Only connected to metal layer 9 power rails (top level local power grid)




Global power rails at chip level

17

Processor Abutment

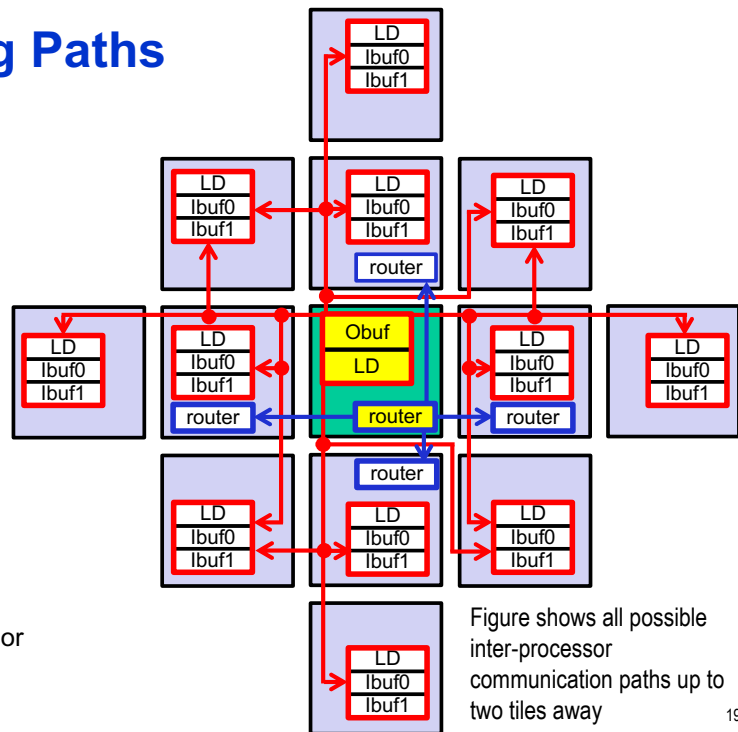
- With 1000 homogeneous processors, very tight abutment is important
 - Communication wires were constrained with I/O pin placement
 - 0.4 μm horizontal spacing and 4.5 μm vertical spacing
 - Required fabrication cells, such as alignment cells cause interruption
 - Removing a processor is non-ideal due to timing
 - Specific rows required added vertical space to fit these cells



18


Processor Timing Paths

- 196 independently timed datapaths for direct communication up to *only two* tile hops
- Possible timing paths
 - **Circuit switched (x2 layers):**
 - Each core has 2 origins:
 - One output buffer (Obuf)
 - One long-distance buffer (LD)
 - Each core has 3 destinations:
 - One long-distance buffer (LD)
 - Input buffer 0 (Ibuf0)
 - Input buffer 1 (Ibuf1)
 - Signals can pass through cores
 - **Packet switched:**
 - Single hop routers on each processor




19

Inter-Processor Communication Wires




Circuit switch circuitry
1% area
(9% including FIFOs)



Packet router circuitry
9% area

20

C4 Routing and Chip I/O




- Premade flip-chip BGA package was used
 - 564 C4 bumps
 - 162 I/O (63 LVDS Pairs)
 - 402 Power (Vdd, Gnd, and VddIO)
 - Non-ideal C4 placement and quantity
- All drivers and ESD clamps were on the periphery
- When all cores are at full speed and 100% activity, only sufficient current for center processors
 - Most applications average 50% activity

21

Oscillator and Processor Clock Tree

- Inverter ring oscillator
 - Inverter chain cells were hand selected
 - Placed and routed before processor to maintain timing
- Clock tree with **5,522** leaves to drive the flip-flop memories

Osc.

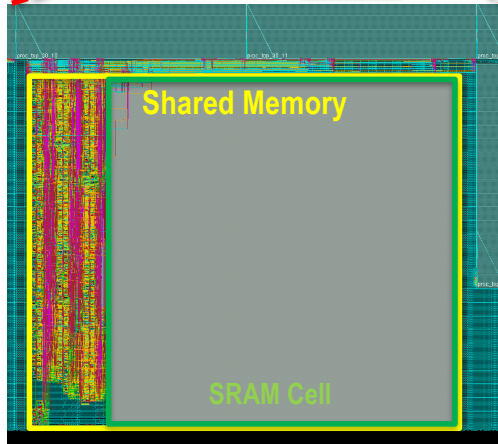



Processor Skew (70 ps max.)

Average Measured Processor Osc. Frequencies	
1.1 V	1.78 GHz
900 mV	1.24 GHz
560 mV	115 MHz


22

Array Placement with Memories

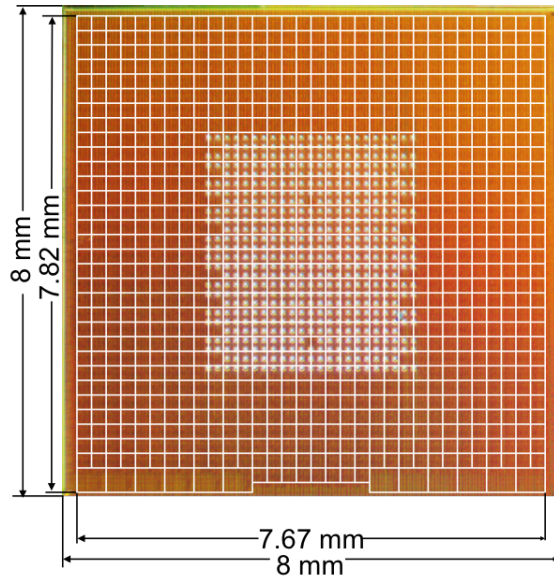


- Embedding memories in the array placement
 - Hard SRAM memory blocks forced the shared memory to take a specific shape
- Memories placed on bottom to eliminate:
 - Configuration signal pass-through
 - Non-regular timing for signals over memory
 - Wasted space

23



KiloCore Chip



Technology	32nm IBM PDSOI CMOS
Num. Procs.	1000
Num. Mem.	12
Die Area	64 mm ²
Array Area	60 mm ²
Transistors	621 Million
C4 Bumps	564 (162 I/O)
Package	676 Pad Flip-Chip BGA

25

Outline

- Introduction and Motivation
 - Core Scaling Trend
- KiloCore Architectural Design
- Physical Chip Design
 - Chip Flow
 - Design Challenges
- Final Die
- **Application Measurements**

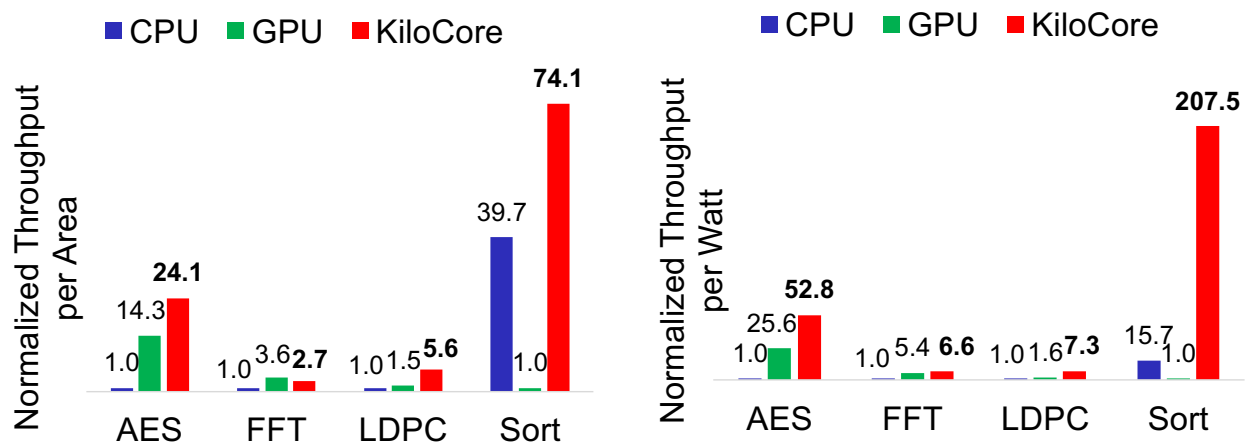
26

Applications

- Several applications have been implemented for KiloCore:
 - Fast Fourier Transform
 - 4096 length, 16-bit fixed-point complex data
 - Advanced Encryption Standard
 - 128-bit keys
 - Low Density Parity Check
 - 4095 code length
 - Sorting
 - 100 Byte records with 10 Byte keys
 - 1850 records per sorted block

27

Application Comparison



- KiloCore operating at 1.1 V

Data: B. Bohnenstiehl *et al.*, JSSC 2017.

28

Acknowledgments

- Funding and Support

- DoD and ARL/ARO Grant W911NF-13-1-0090
- TAPO
- NSF CAREER award 546907
 - CCF Grant No. 430090
 - CCF Grant No. 903549
 - CCF Grant No. 1018972
 - CCF Grant No. 1321163
- SRC GRC Grant 1598
 - CSR Grant 1659
 - GRC Grant 1971
 - GRC Grant 2321
- ST Microelectronics
- C2S2
- Intel Corporation
- UCD Faculty Research Grant
- MOSIS
- Artisan