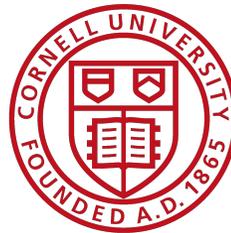


DRAF: A Low-Power DRAM-based Reconfigurable Acceleration Fabric

Mingyu Gao, Christina Delimitrou, Dimin Niu, Krishna Malladi, Hongzhong Zheng,
Bob Brennan, Christos Kozyrakis



FPGA-Based Accelerators

- ❑ Improve performance and energy efficiency
- ❑ Good balance between flexibility (CPUs) and efficiency (ASICs)

- ❑ Recently used for many datacenter apps
 - Image/video processing, websearch, neural networks, ...



Motivation

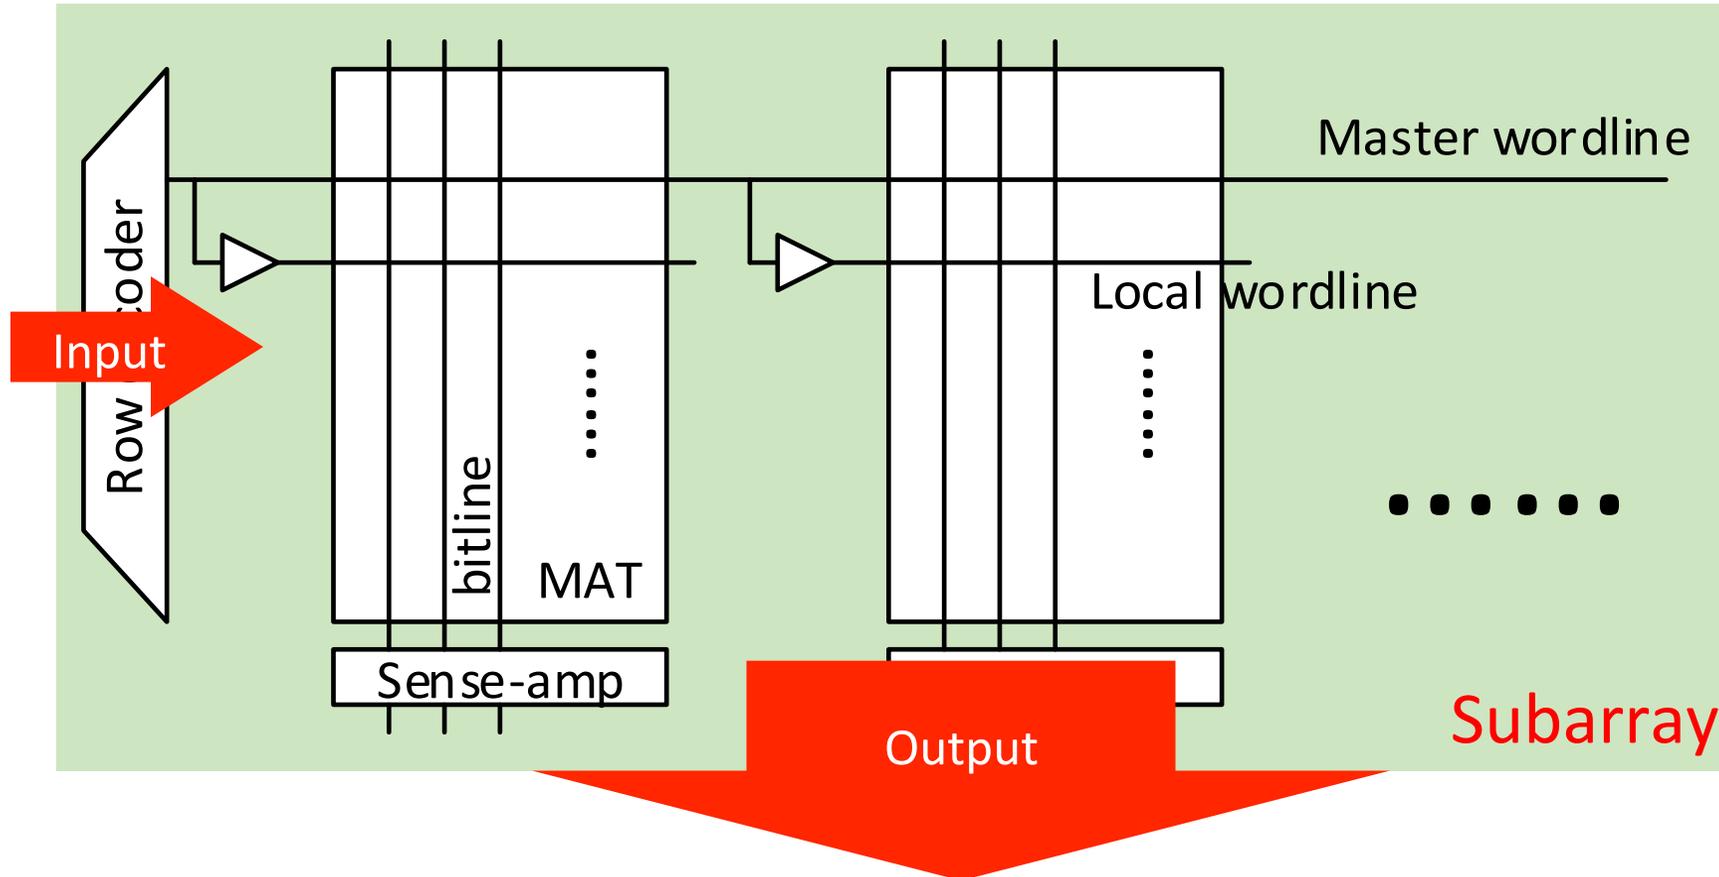
- ❑ Deploy FPGAs in cost & power constrained systems
- ❑ Datacenter systems
 - *High-density* FPGAs for large accelerators for multiple apps
 - *Low-power* FPGAs to simplify integration in servers and racks
- ❑ Mobile systems
 - *High-density* FPGAs for accelerators for multiple apps
 - *Low-power* FPGAs for low cost and long battery life

DRAF in a Nutshell

- ❑ A *high-density & low-power* FPGA
 - Bit-level reconfigurable, just like conventional FPGAs
- ❑ Uses dense *DRAM technology* for lookup tables
 - Replacing the SRAM technology in conventional FPGAs
- ❑ DRAF vs. FPGA
 - 10 – 100x logic density
 - 1/3 power consumption
 - Multi-context support with fast context switch

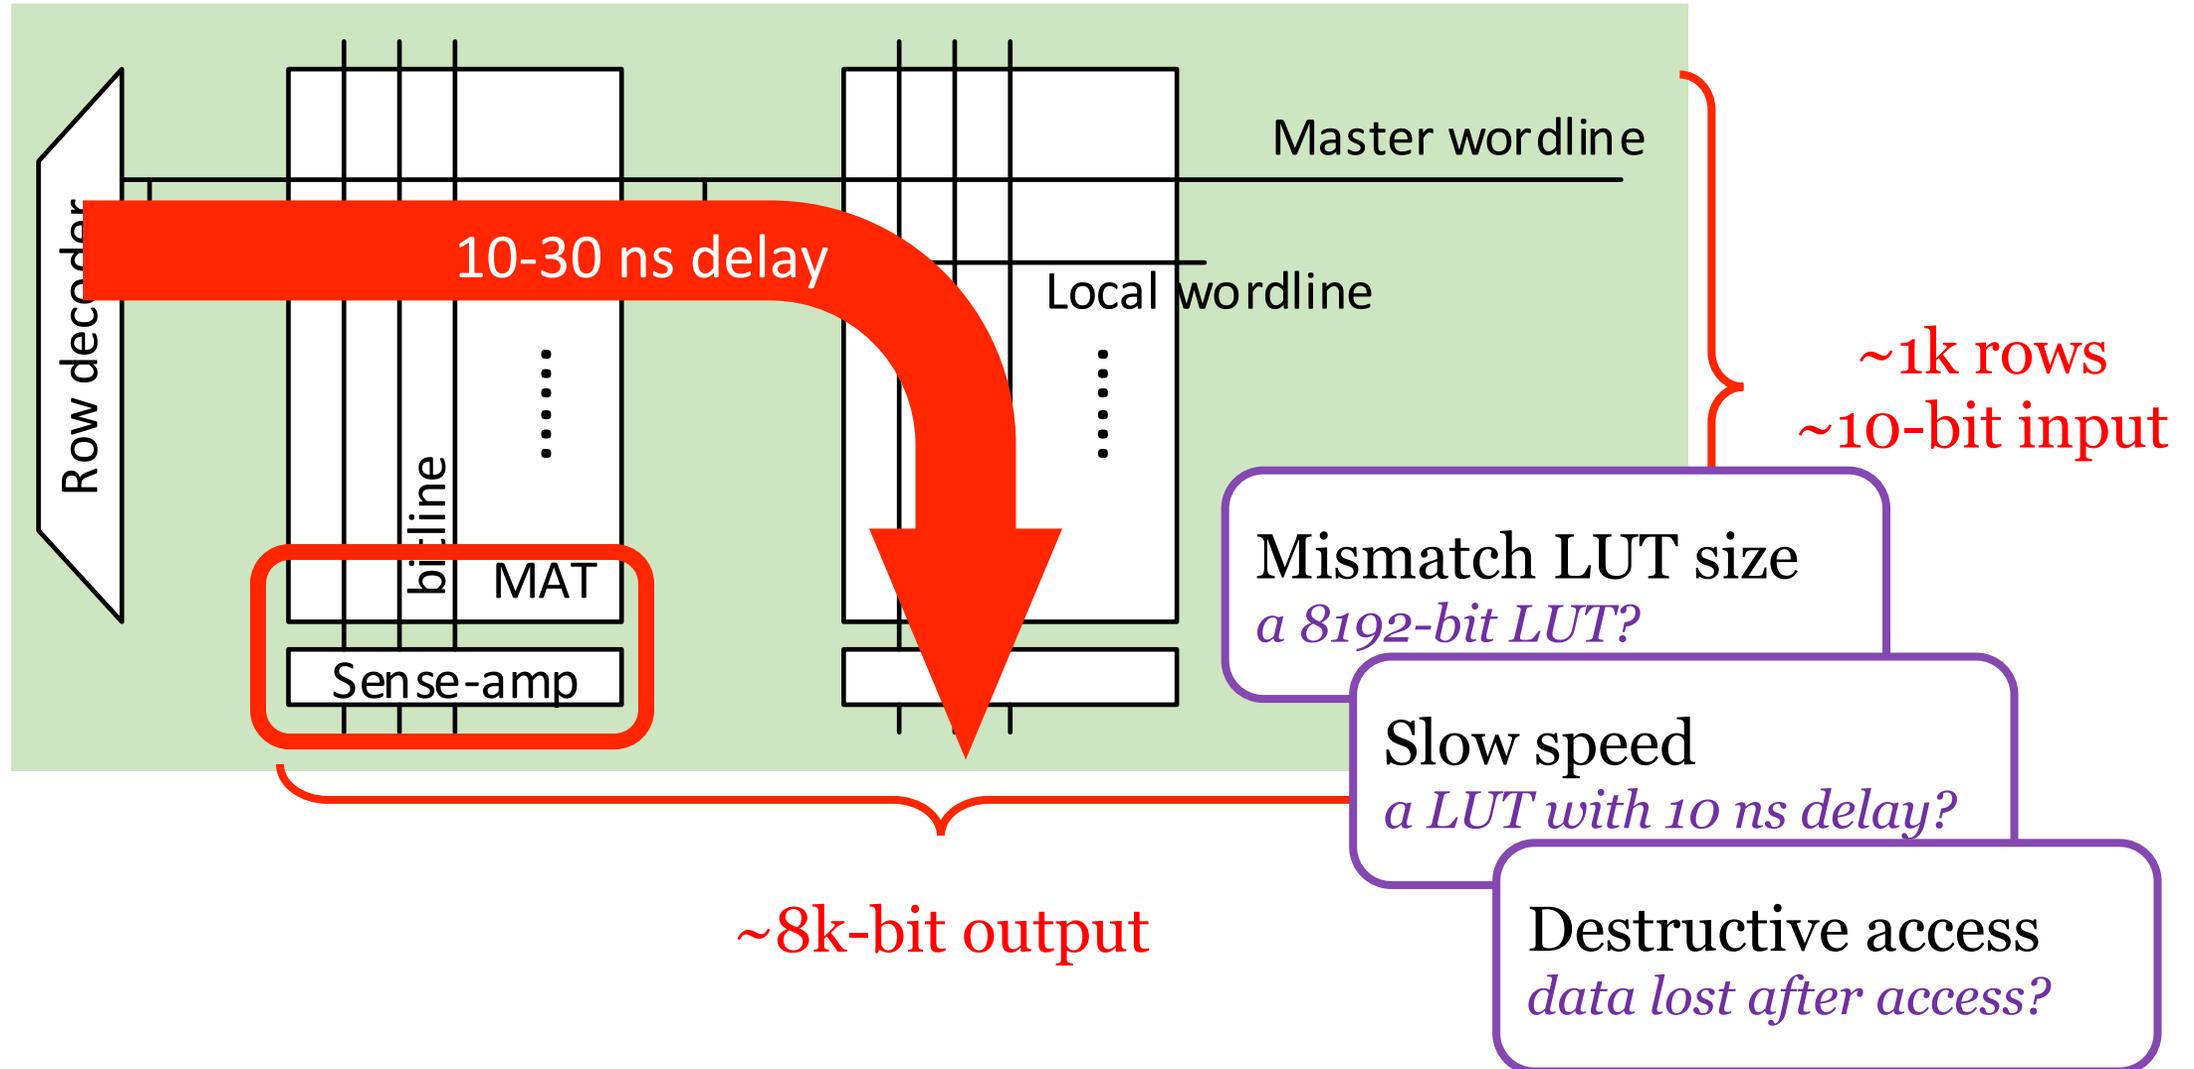
Challenges of Building DRAM-based FPGAs

DRAM Array Structure



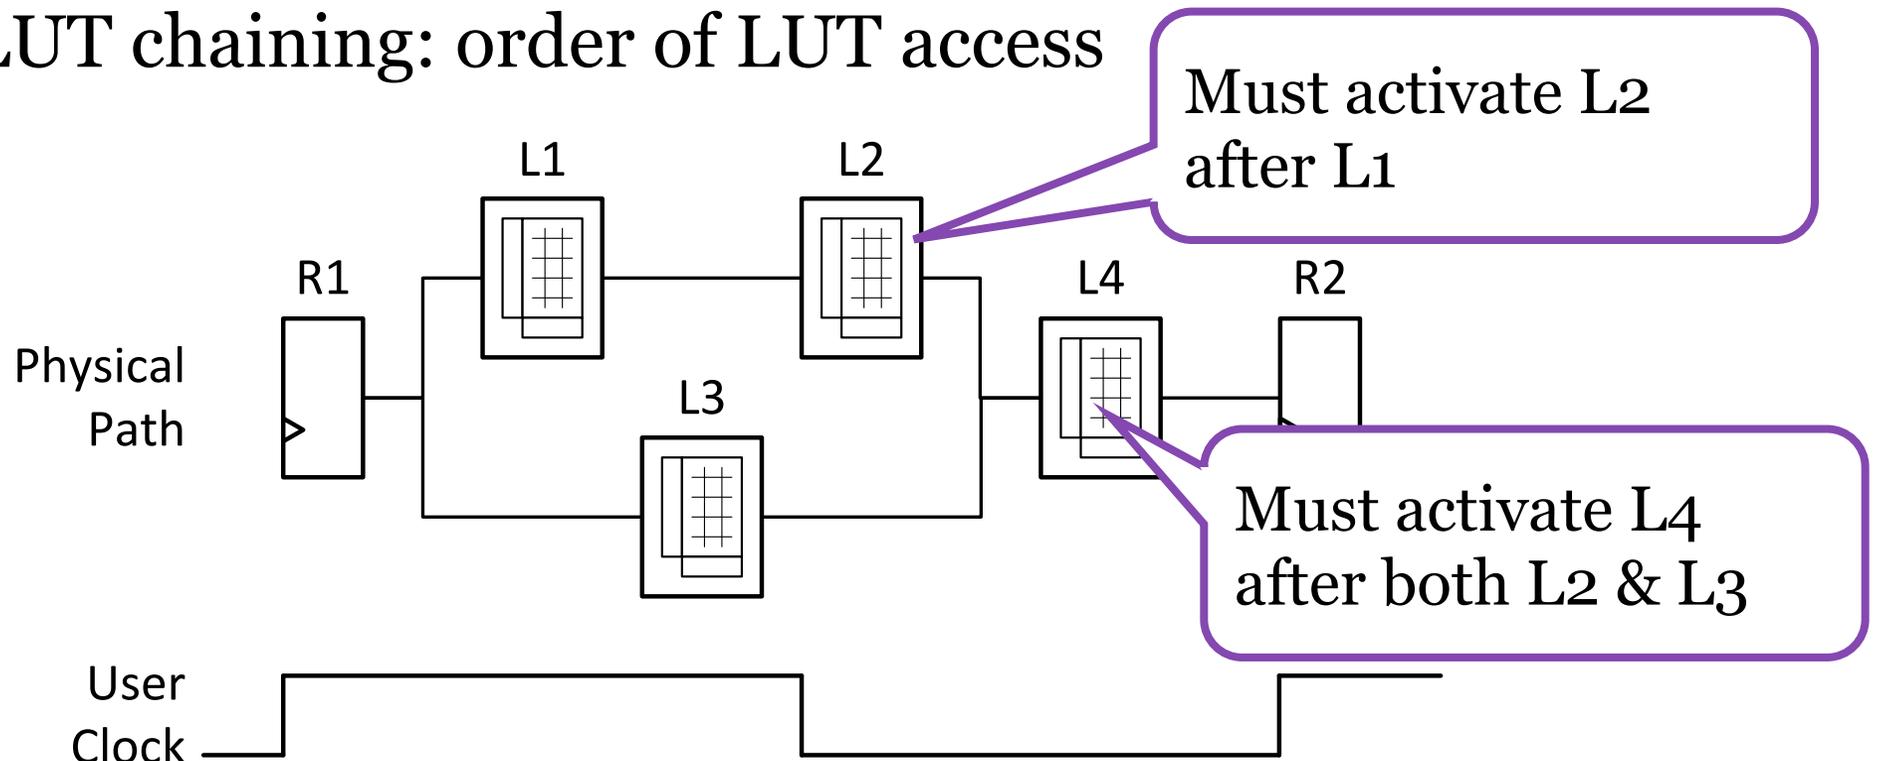
A DRAM subarray is naturally a lookup-table

Challenges



Destructive Access

- ❑ Explicit activation, restoration, and precharge operations
 - Longer access delay due to serialization
- ❑ Issue of LUT chaining: order of LUT access



DRAF Architecture

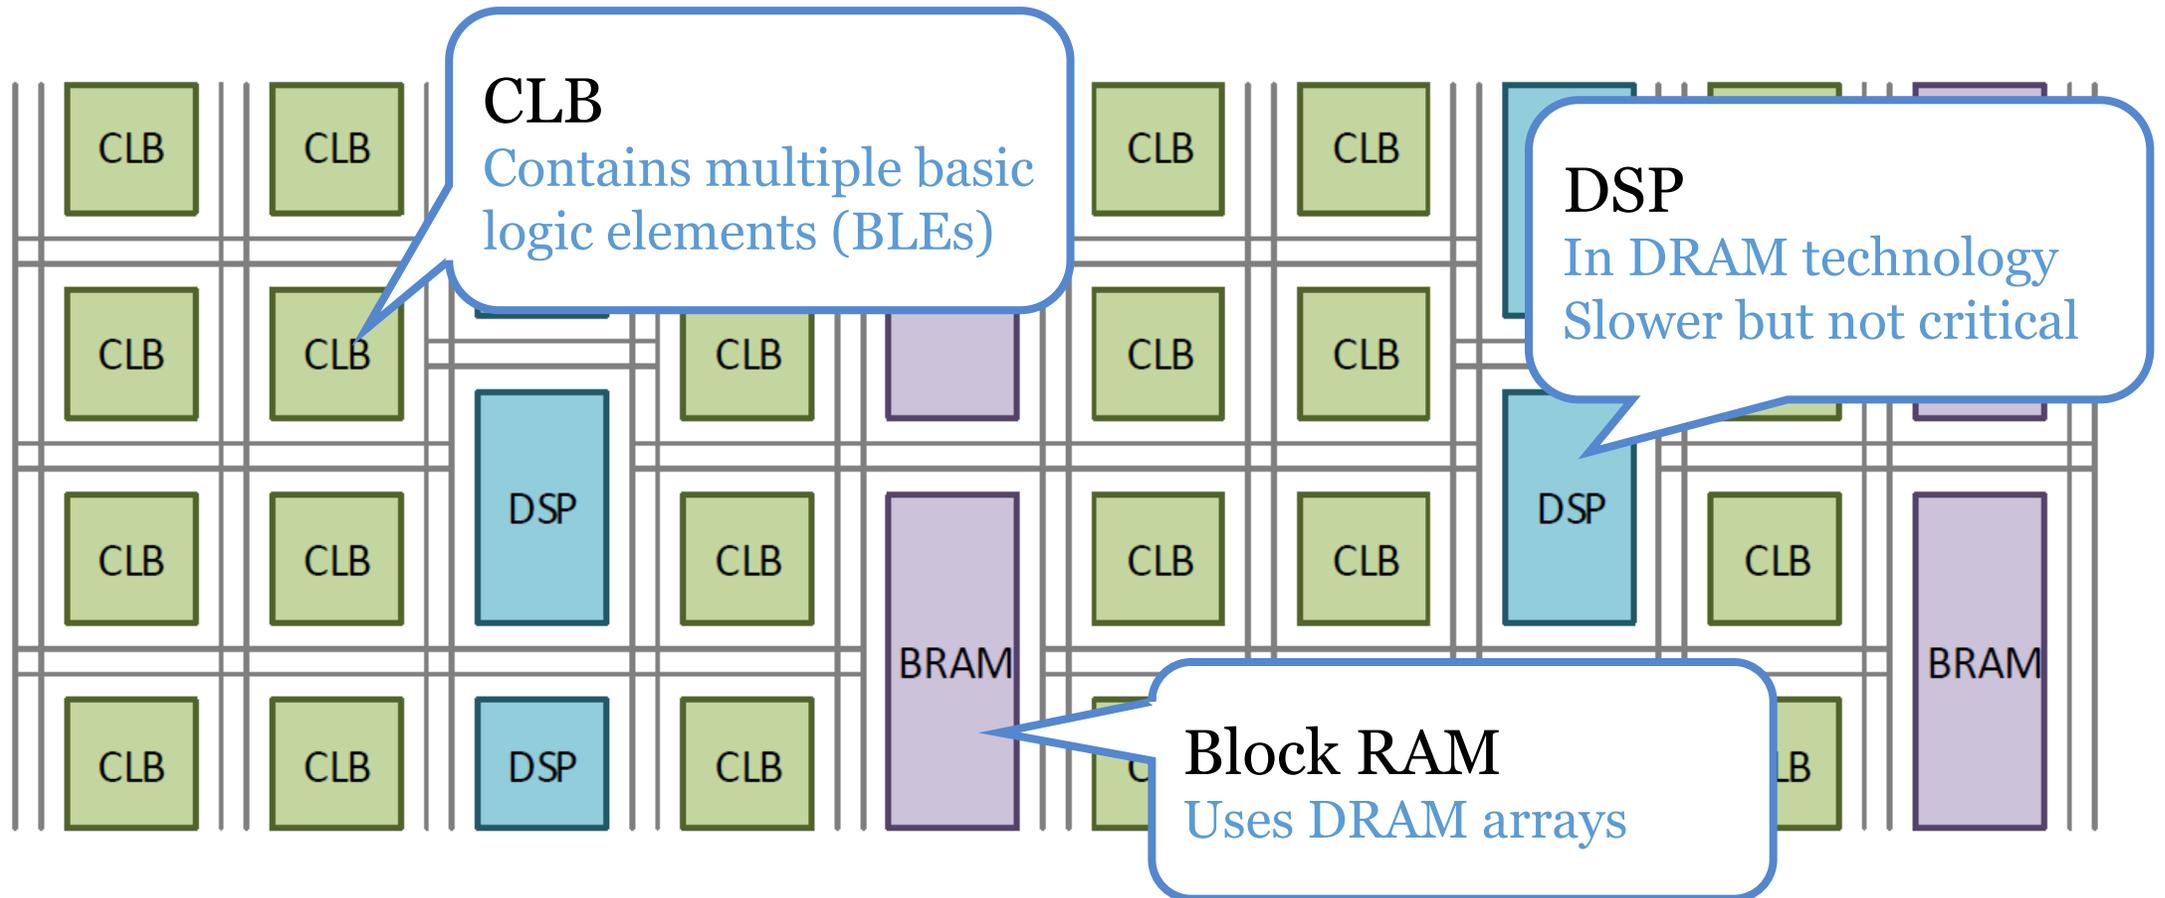
Basic Logic Element

Multi-Context Support

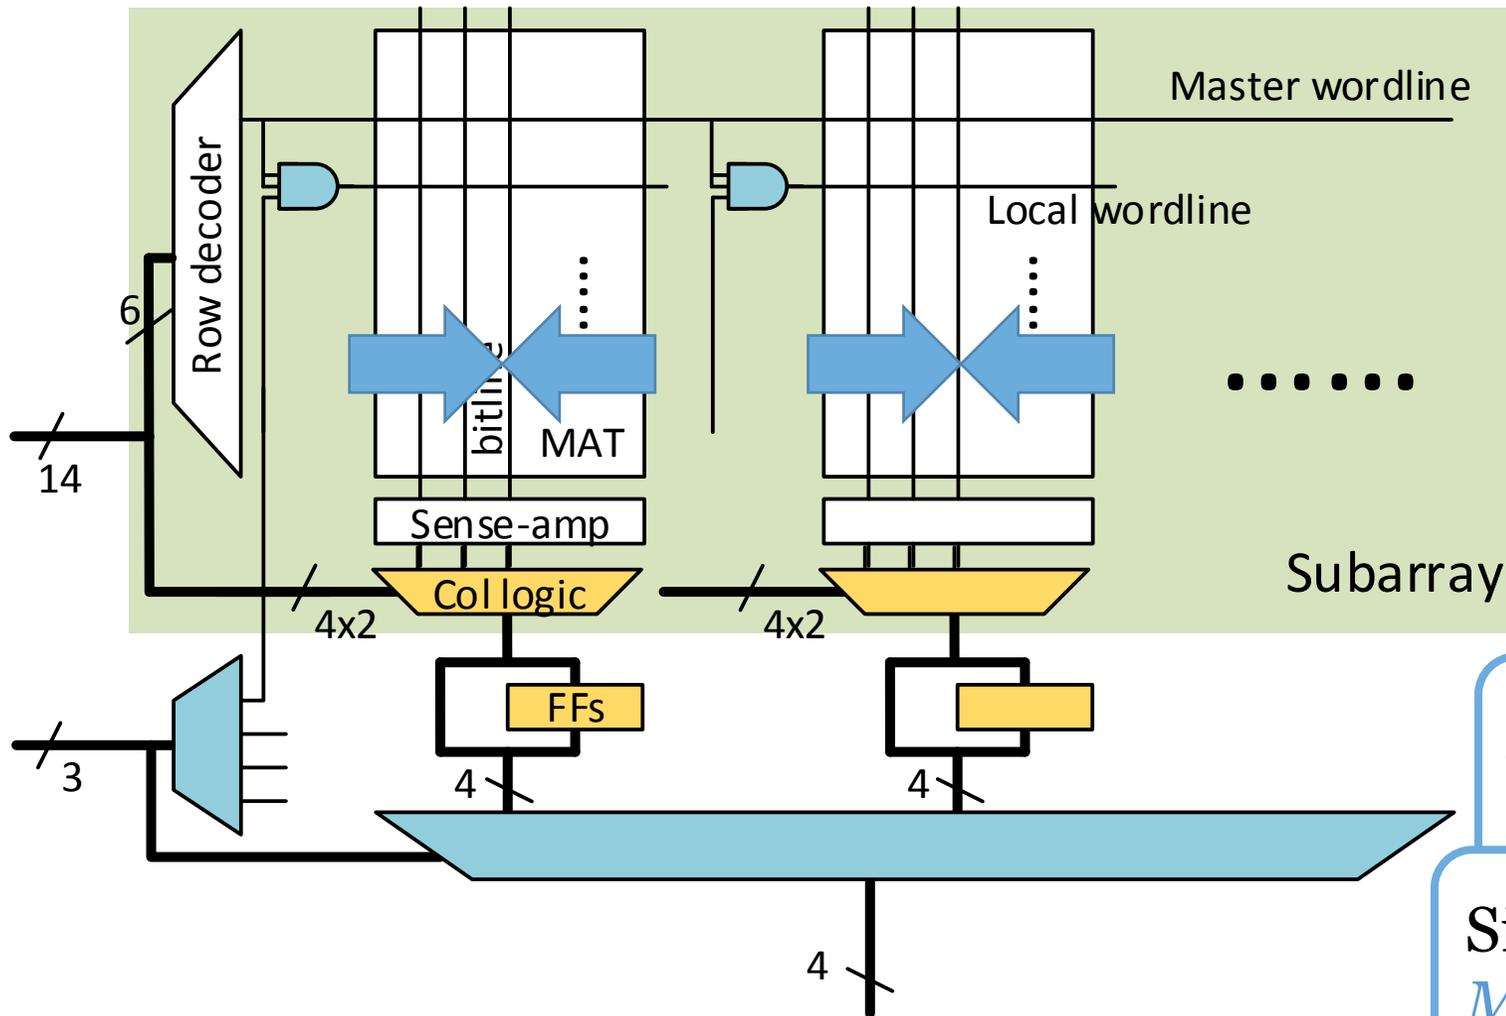
Timing

DRAF Overview

- Same island layout and configurable interconnect as FPGA



Basic Logic Element



7-10 bits input
2-4 bits output

Narrower MAT
1k bits to 8-16 bits

Specialized column logic
Better flexibility

Additional FFs & MUXs
Registering & retiming

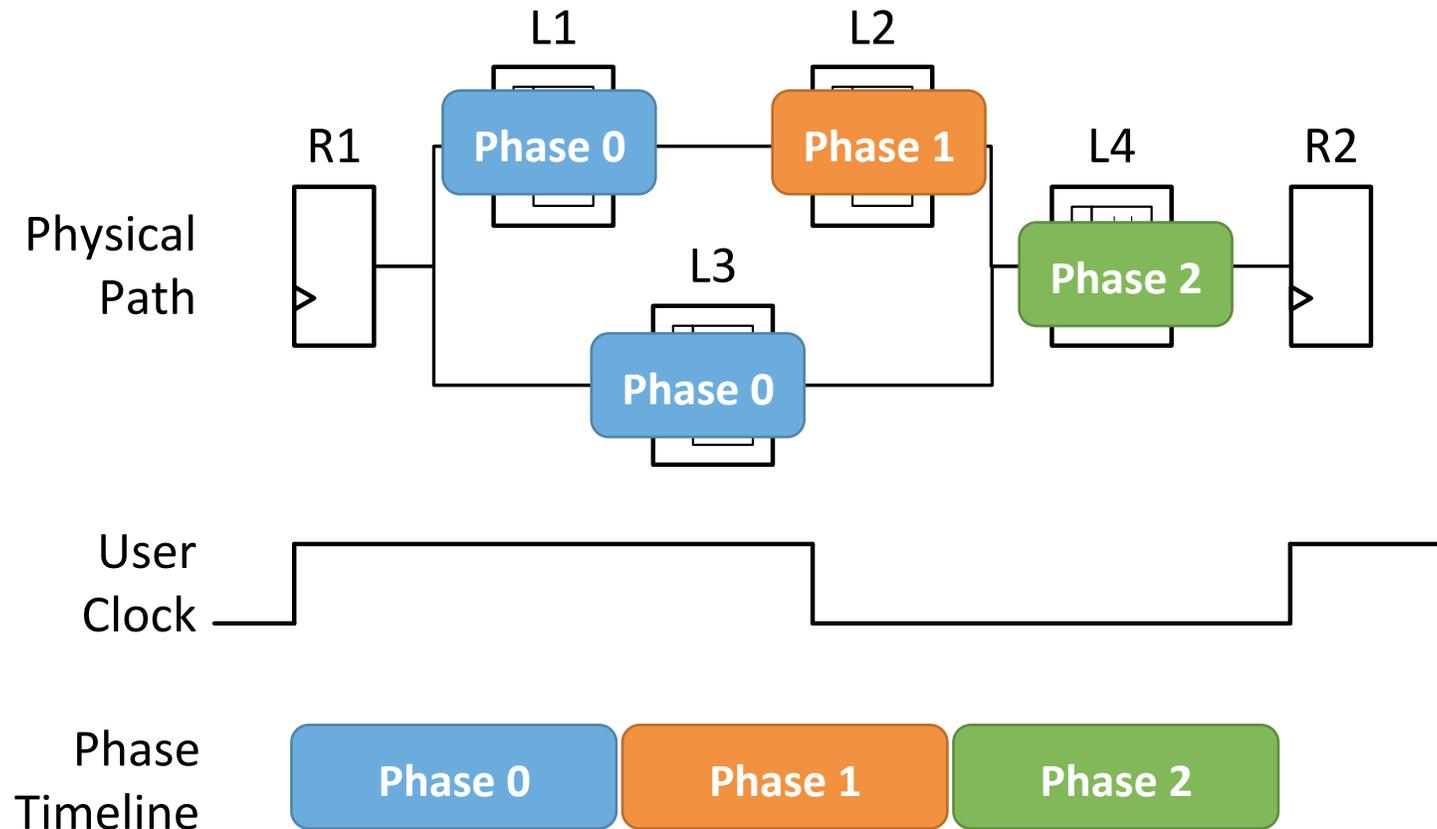
Single-MAT access
Multi-context

Multi-Context Support

- ❑ DRAF supports 8-16 contexts per chip
 - Context: one MAT per BLE
 - Efficient use of MATs with little area and power overhead
- ❑ Instant switch between active contexts
 - Similar to context-switch between processes on CPU
- ❑ Context uses
 - One context per accelerator design or application
 - One context per part of a very large accelerator design

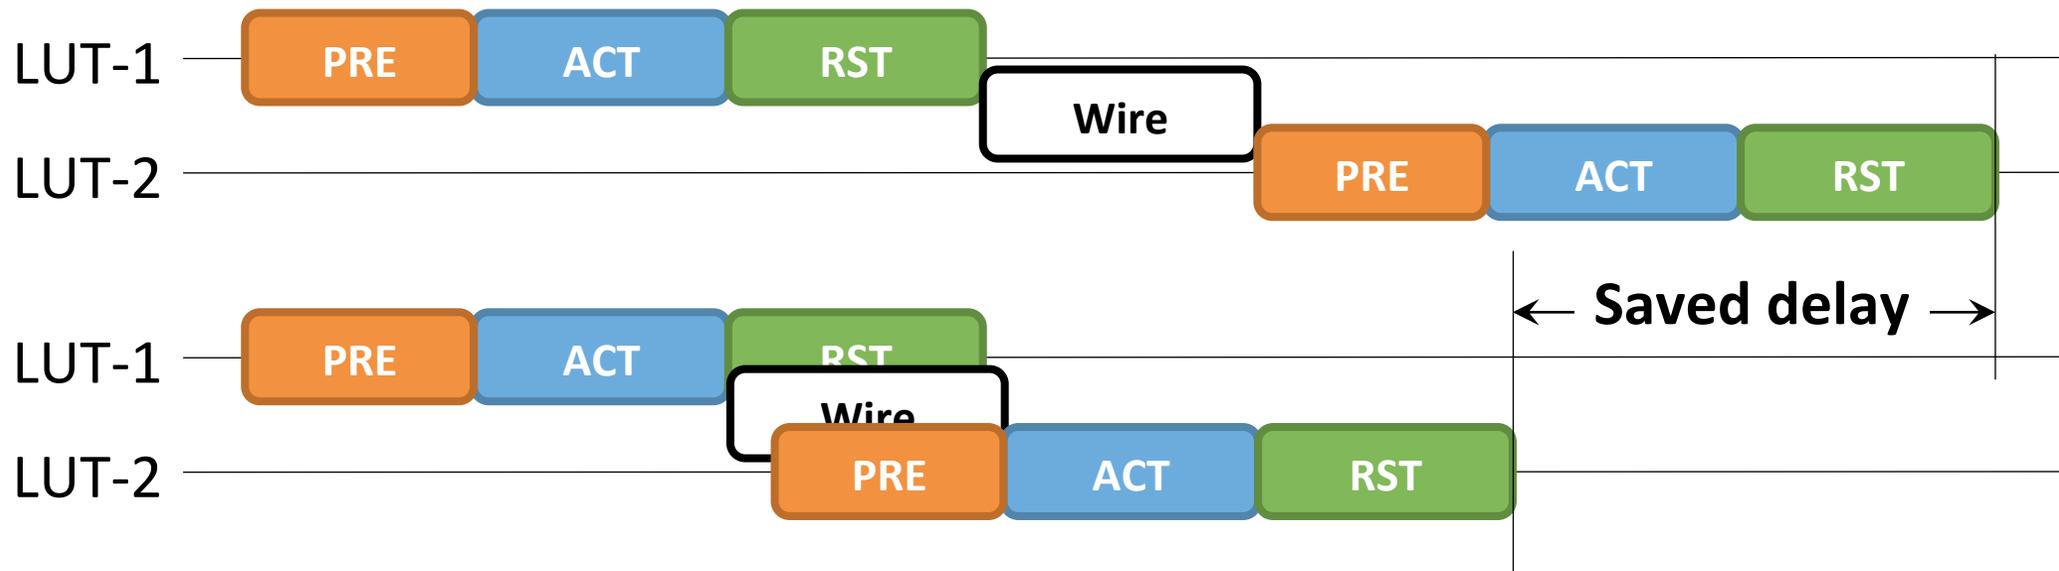
Timing – Destructive Access

- ❑ Issue of LUT chaining: order of LUT access
- ❑ Solution: *phase* – similar to critical path finding



Timing – Latency Optimization

- ❑ Issue: precharge and restore delays
- ❑ Solution: *3-way delay overlapping*
 - Hide PRE/RST delays with wire propagation delay
- ❑ Performance gap between DRAF and FPGA reduces from >10x to 2-4x



Summary

- ❑ Challenges → solutions
 - Mismatch LUT size → *multi-context BLE*
 - Destructive access → *phase-based timing*
 - Slow speed → *3-way delay overlapping*

- ❑ Other design features (see paper)
 - Sense-amp as register
 - Time-multiplexed routing
 - Handling DRAM Refresh

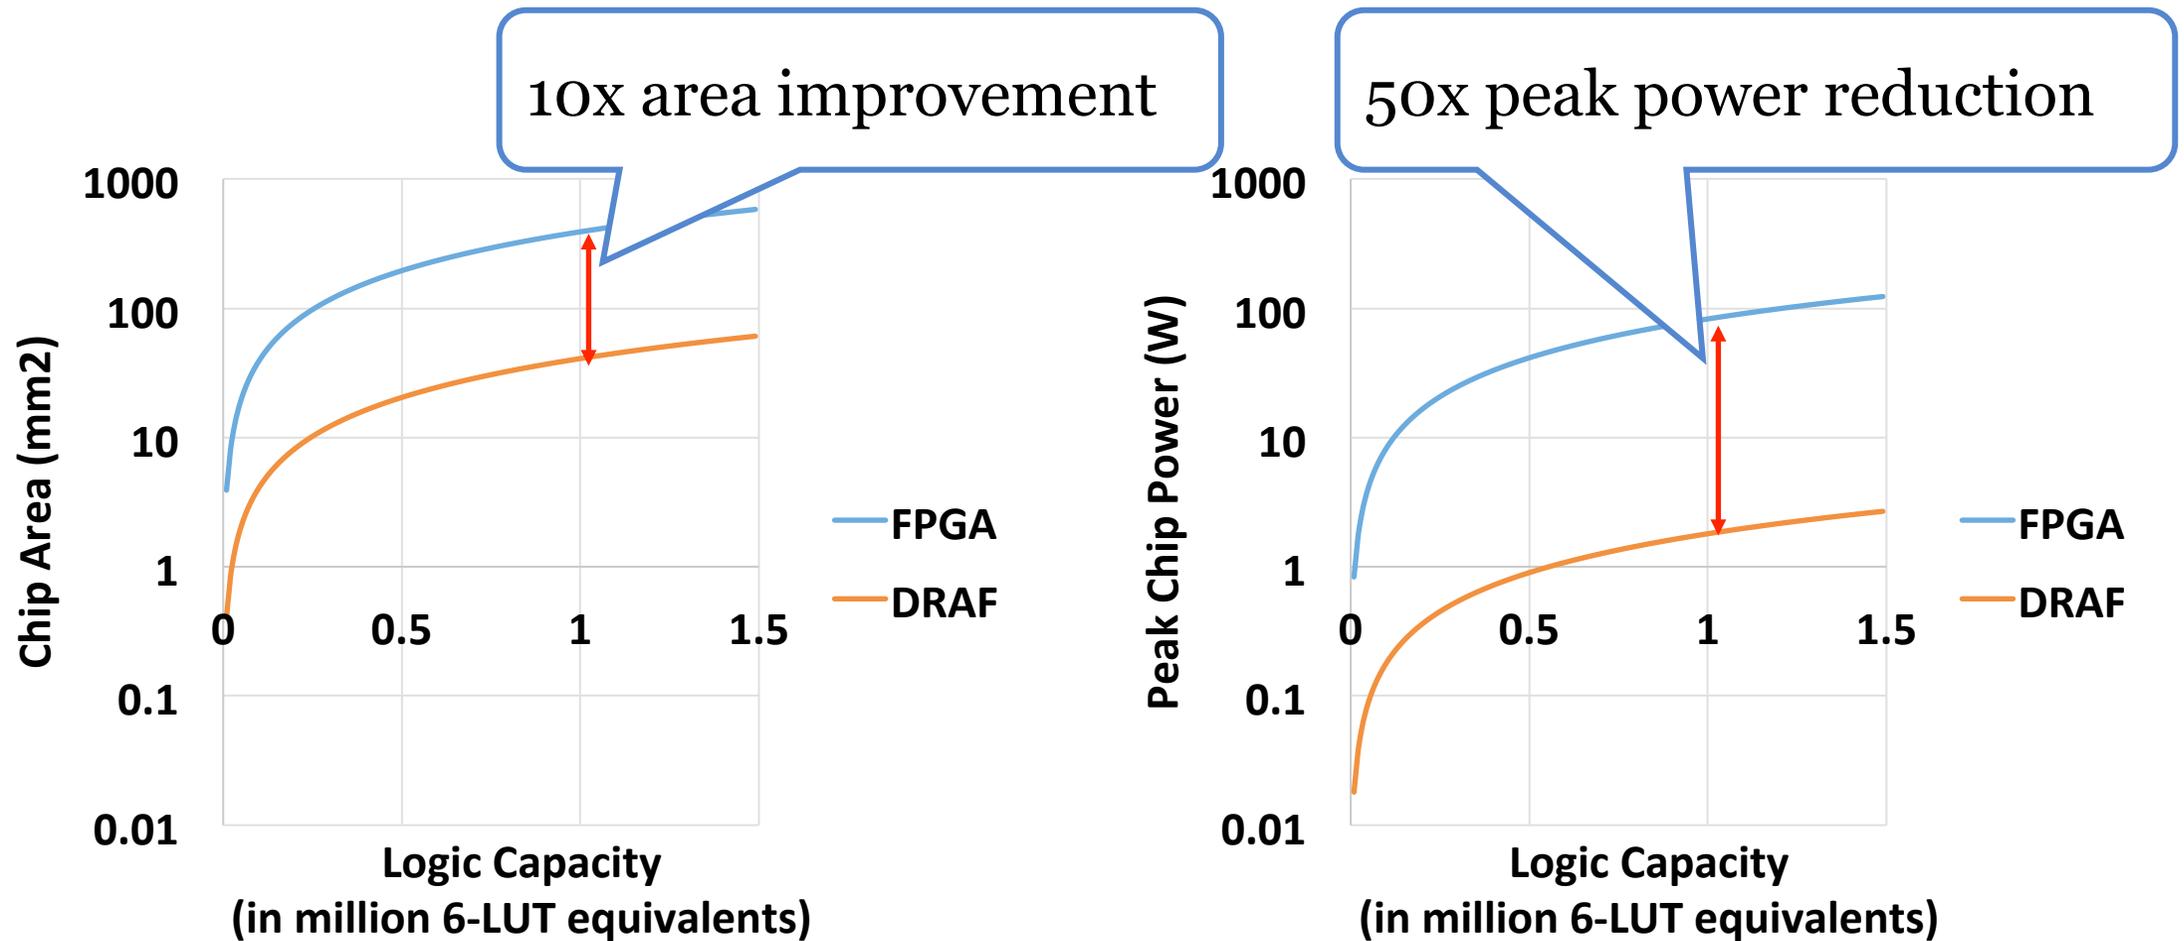
Evaluation

Area, power, performance against FPGA and CPU

Methodology

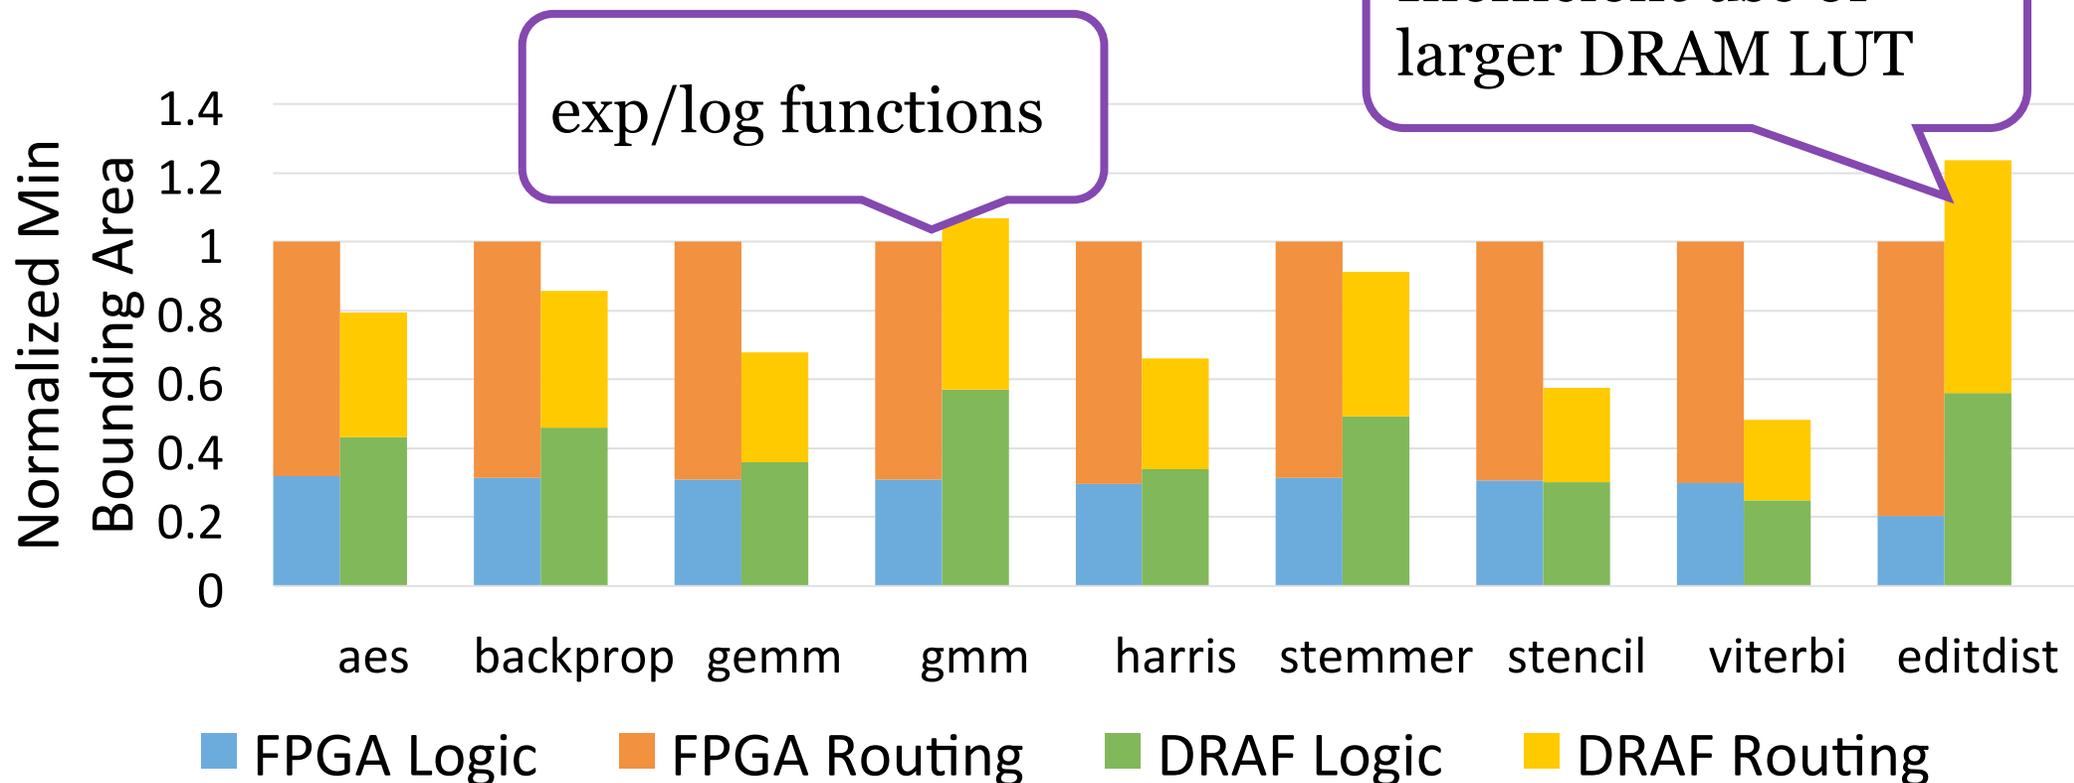
- ❑ Synthesize, place & route with Yosys + VTR
- ❑ CACTI-3DD with 45 nm power and area models
- ❑ Comparisons
 - 70 mm² FPGA based on Xilinx Virtex-6
 - 70 mm² DRAF device, 8-context
 - Intel Xeon E5-2630 multi-core processor (2.3 GHz)
- ❑ 18 accelerator designs
 - MachSuite, Sirius, Vivado HLS Video Library, VTR benchsuite
 - Web service, image processing, analytics, neural networks, ...

DRAF Chip Area & Power



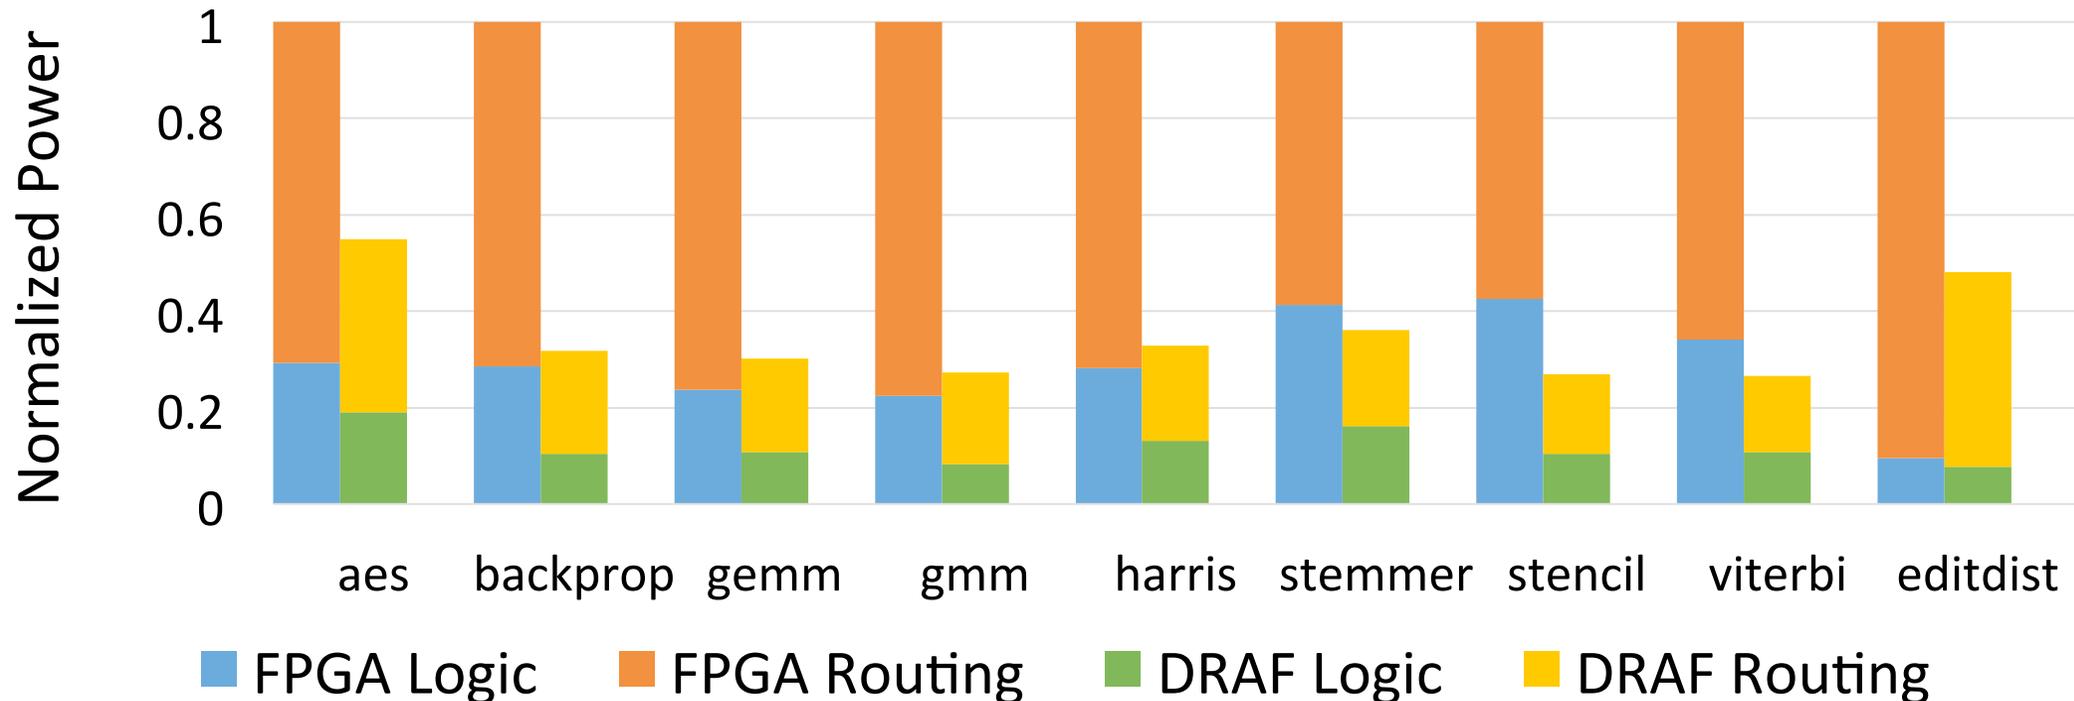
FPGA vs. DRAF (Area)

- ❑ 8-context DRAF occupies 19% less area than 1-context FPGA
 - 10x area efficiency: 8 designs in less silicon area than 1 design before
 - But only one context can be active at a time



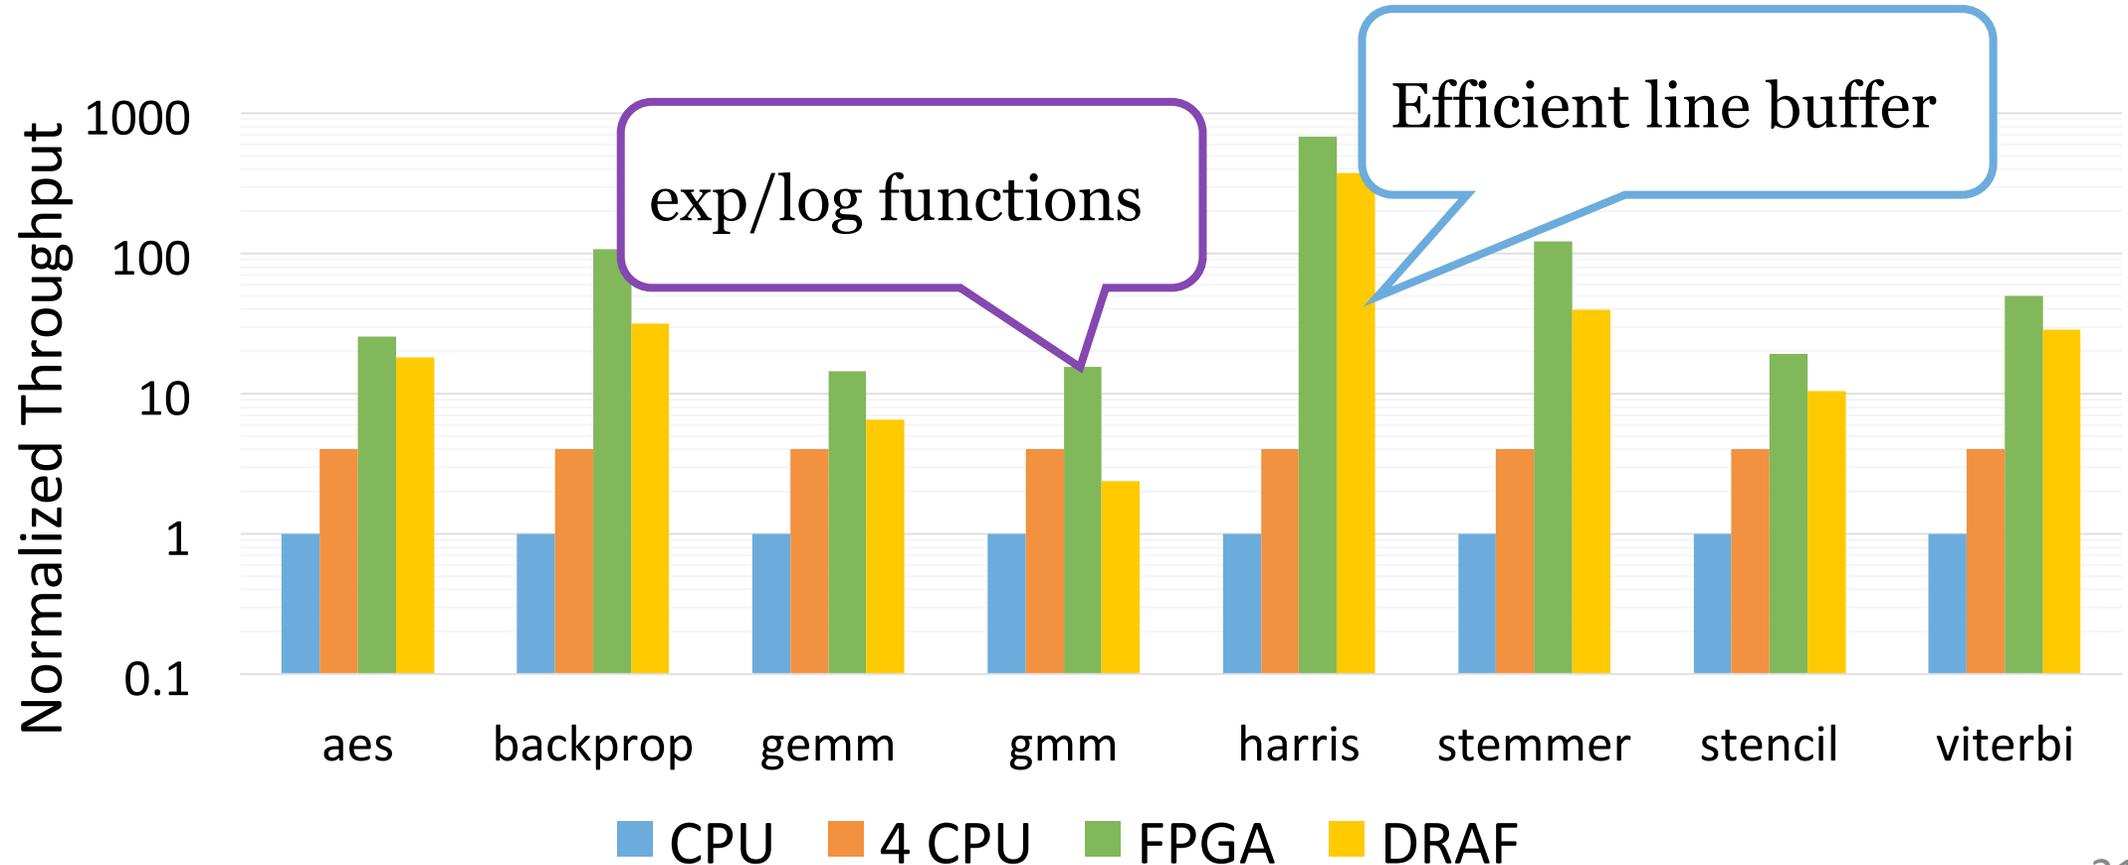
FPGA vs. DRAF (Power)

- ❑ Use one context in DRAF
- ❑ DRAF consumes 1/3 power of FPGA and 15% less energy
 - Note: current CAD tools are less efficient with DRAF



Performance

- ❑ DRAF is 2.7x slower than FPGA
- ❑ DRAF is 13.5x faster than CPU, 3.4x faster than ideal 4-core



Conclusions

- ❑ DRAF: high-density and low-power reconfigurable fabric
 - Based on dense DRAM technology
 - Optimized timing + multi-context support
- ❑ DRAF targets cost and power constrained applications
 - E.g., datacenters and mobile systems
- ❑ DRAF trades off some performance for area & power efficiency
 - 10x smaller area, 3x less power, and 2.7x slower than FPGA
 - Still 13x speedup over Xeon cores

Thanks!

Questions?



Stanford MAST



**Memory
Solutions
Lab**

Backup

Design flow

- ❑ Verilog/VHDL programming and similar synthesis flow
 - DRAF has the same primitives (LUT, FF, DSP, BRAM) as FPGA
- ❑ Specific tweaks
 - Wider LUT: more efficient packing
 - Optimize for latency rather than area
 - Routing delay is easier to handle
 - Additional timing requirements, e.g. phase, etc.

Multi-Context

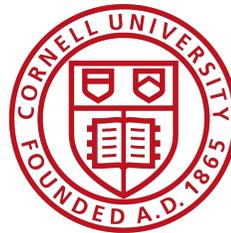
- ❑ Why not do multi-context in SRAM FPGAs?
- ❑ Store contexts in-place
 - High area overhead, can be use to implement more normal LUTs
 - In DRAF: little overhead due to dense DRAM MAT array
- ❑ On-chip backup storage
 - Significant context switch overheads in power and latency
 - In DRAF: *zero* latency and power for context switch

Design Exploration

- ❑ Lots of data in paper
- ❑ Main tradeoff is between area and latency
 - Larger LUT: better area, worse latency
 - Smaller LUT: worse area, better latency
- ❑ A major limitation is the CAD tool
 - Cannot efficiently map applications to large LUTs
- ❑ Final LUT size
 - 7-input, 2-output, 8-context
 - 64 rows, 32 columns, 2048-bit subarray

DRAF: A Low-Power DRAM-based Reconfigurable Acceleration Fabric

Mingyu Gao, Christina Delimitrou, Dimin Niu, Krishna Malladi, Hongzhong Zheng,
Bob Brennan, Christos Kozyrakis



The Need for High-density & Low-power FPGAs

- ❑ FPGA accelerators improve performance and energy efficiency
 - Recently used for many datacenter apps (Microsoft, Baidu, ...)
- ❑ Datacenter systems
 - Need *high-density* FPGAs for large accelerators for multiple apps
 - Need *low-power* FPGAs to simplify integration in servers and racks
- ❑ Mobile systems
 - Need *high-density* FPGAs for accelerators for multiple apps
 - Need *low-power* FPGAs for low cost and long battery life

DRAF: A High-density & Low-power FPGA

- ❑ Based on *dense DRAM arrays* instead of SRAM LUTs
 - 10-100x density of conventional FPGAs
 - 1/3 power consumption of conventional FPGAs
 - 13x speedup over Xeon cores
- ❑ Come to the talk to learn about
 - Dense, slow DRAM arrays as small, fast LUTs
 - Phase-based timing to address the problem of destructive reads
 - Multi-context support with instantaneous context switch
- ❑ Session 8A, Wednesday 9am