

A Fast Large-Integer Extended GCD Algorithm and Hardware Design for Verifiable Delay Functions and Modular Inversion

Kavya Sreedhar, Mark Horowitz, Christopher Torng

Stanford University

skavya@stanford.edu

January 18, 2023

Extended GCD

Computes Bézout coefficients $\mathbf{b}_a, \mathbf{b}_b$ satisfying Bézout's Identity

$$\mathbf{b}_a, \mathbf{b}_b : \mathbf{b}_a * a_0 + \mathbf{b}_b * b_0 = \text{gcd}(a_0, b_0)$$

Extended GCD is widely used in cryptography

Computes Bézout coefficients $\mathbf{b}_a, \mathbf{b}_b$ satisfying Bézout's Identity

$$\mathbf{b}_a, \mathbf{b}_b : \mathbf{b}_a * a_0 + \mathbf{b}_b * b_0 = \text{gcd}(a_0, b_0)$$

Modular Multiplicative Inverse

RSA

Elliptic Curve Cryptography

ElGamal Encryption

⋮

There is an increasing need for faster XGCD

1. Modular Inversion for Curve25519 [Ber06]
 - Constant-time XGCD faster than Fermat's Little Theorem [BY19]

$$x^{-1} = x^{p-2} \pmod{p}$$

There is an increasing need for faster XGCD

1. Modular Inversion for Curve25519 [Ber06]
 - Constant-time XGCD faster than Fermat's Little Theorem [BY19]

$$x^{-1} = x^{p-2} \pmod{p}$$

2. Squaring binary quadratic forms over class groups [Wes19] as a VDF
 - XGCD is the bottleneck [BBBF18]

$$f(x) = x^{2^T} \text{ in a class group}$$

There is an increasing need for faster XGCD

1. Modular Inversion for Curve25519 [Ber06]
 - Constant-time XGCD faster than Fermat's Little Theorem [BY19]

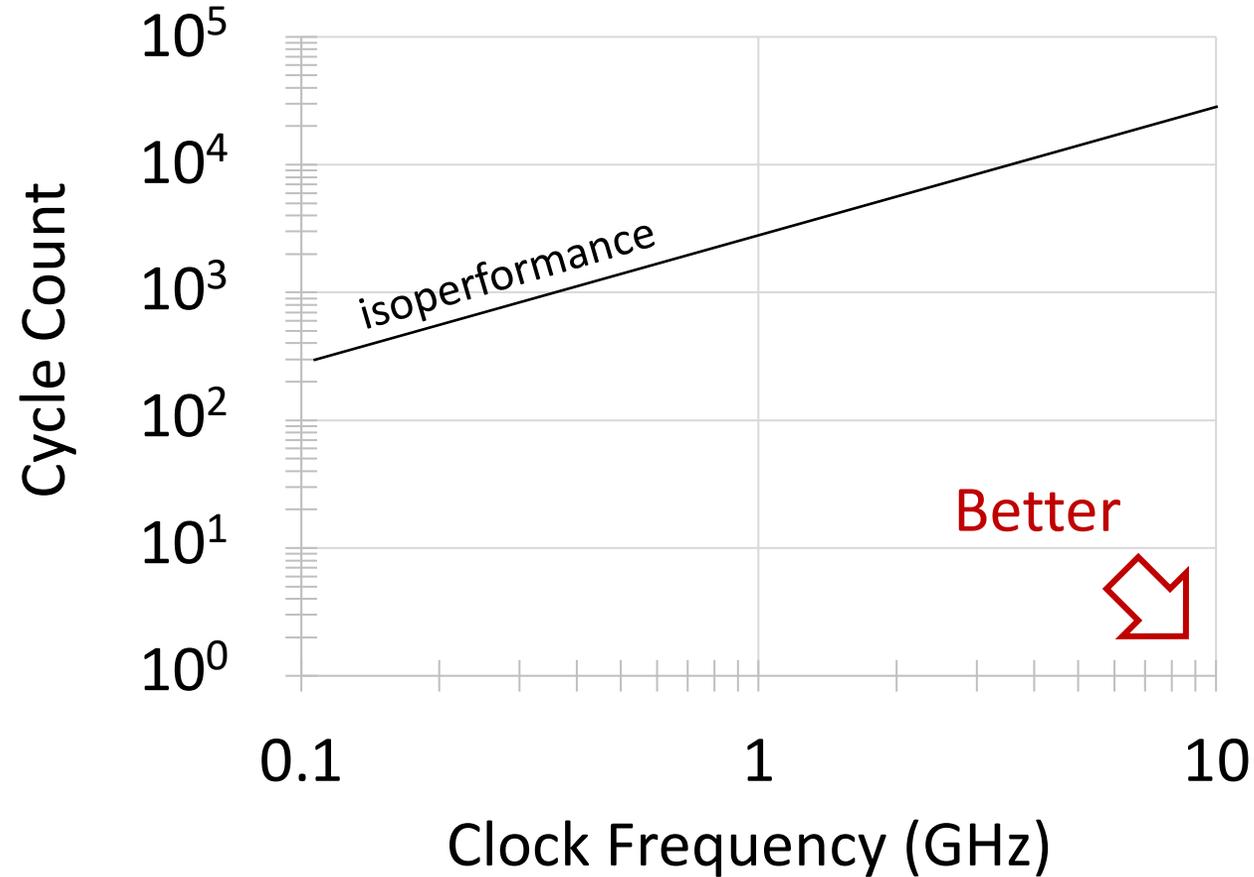
255-bits, constant-time

$$x^{-1} = x^{p-2} \pmod{p}$$

2. Squaring binary quadratic forms over class groups [Wes19] as a VDF
 - XGCD is the bottleneck [BBBF18]

1024-bits, not constant-time

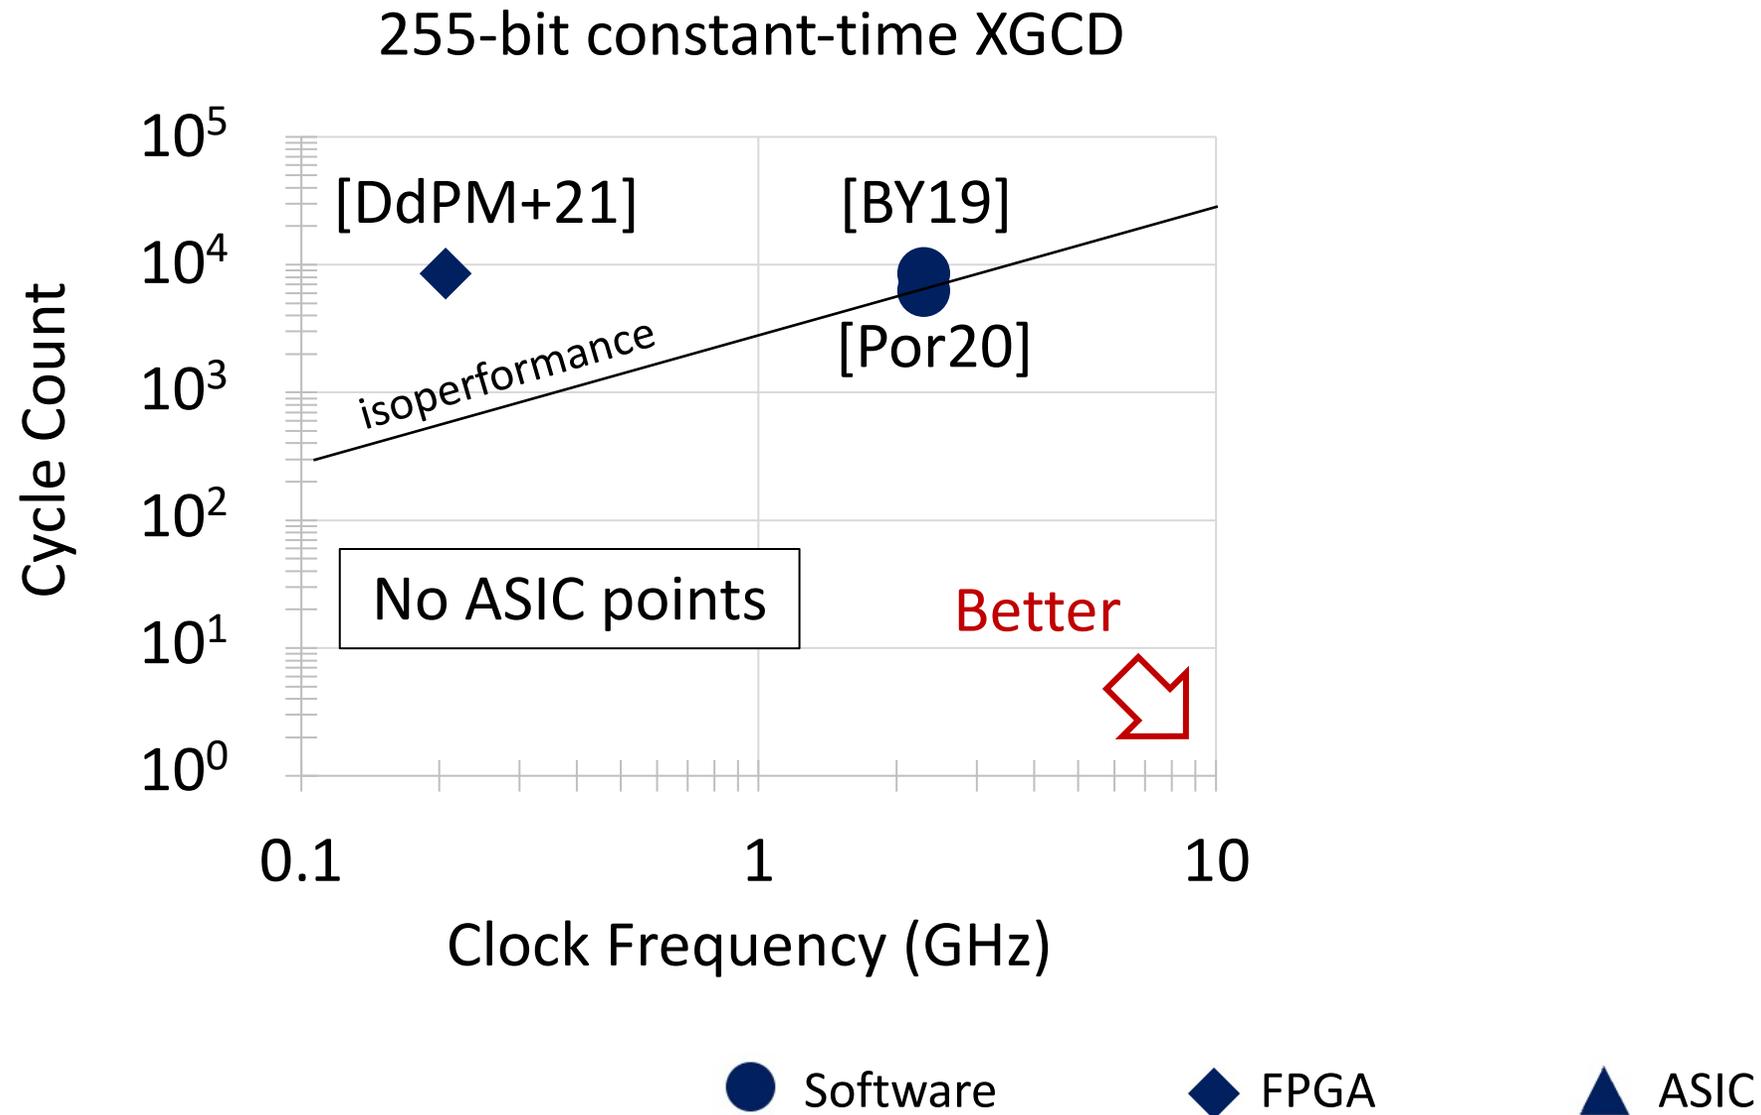
$$f(x) = x^{2^T} \text{ in a class group}$$



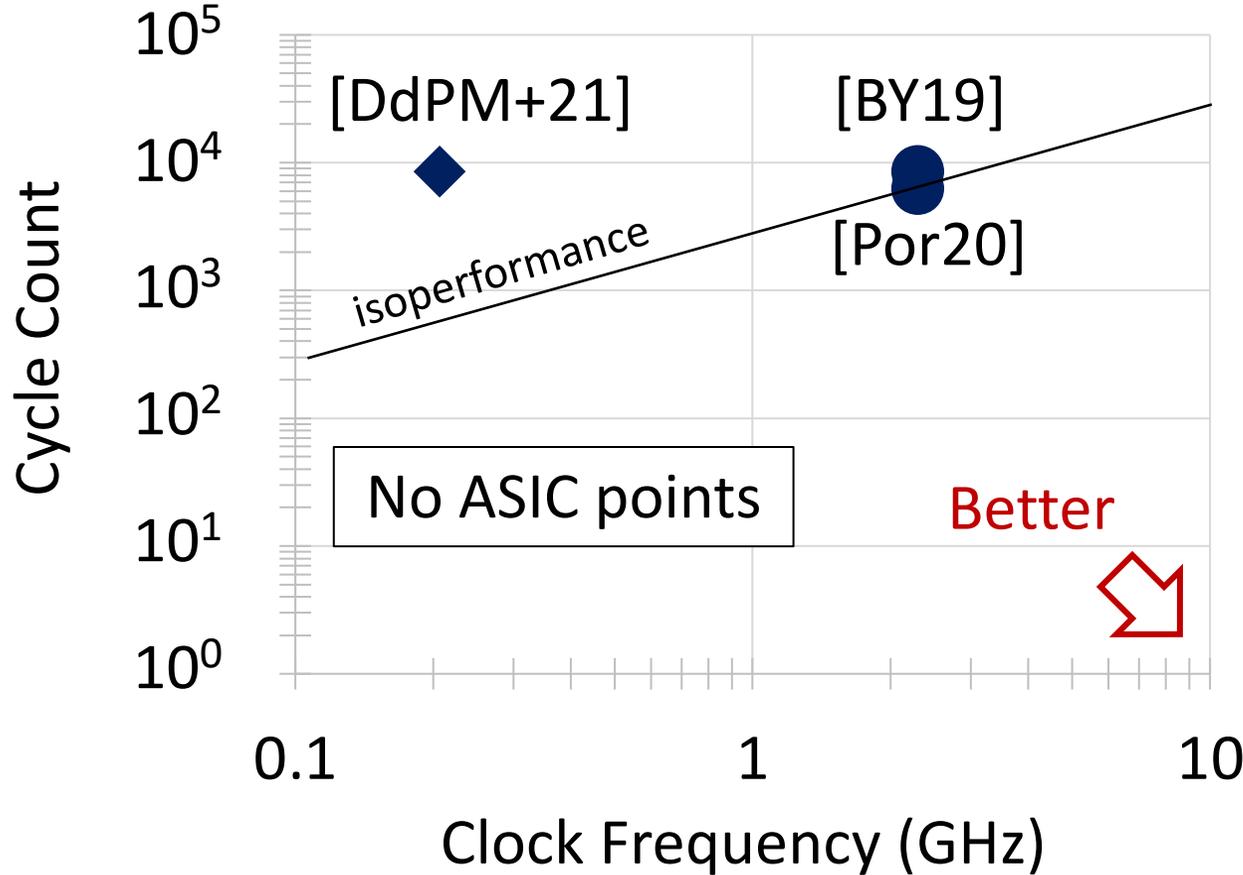
● Software

◆ FPGA

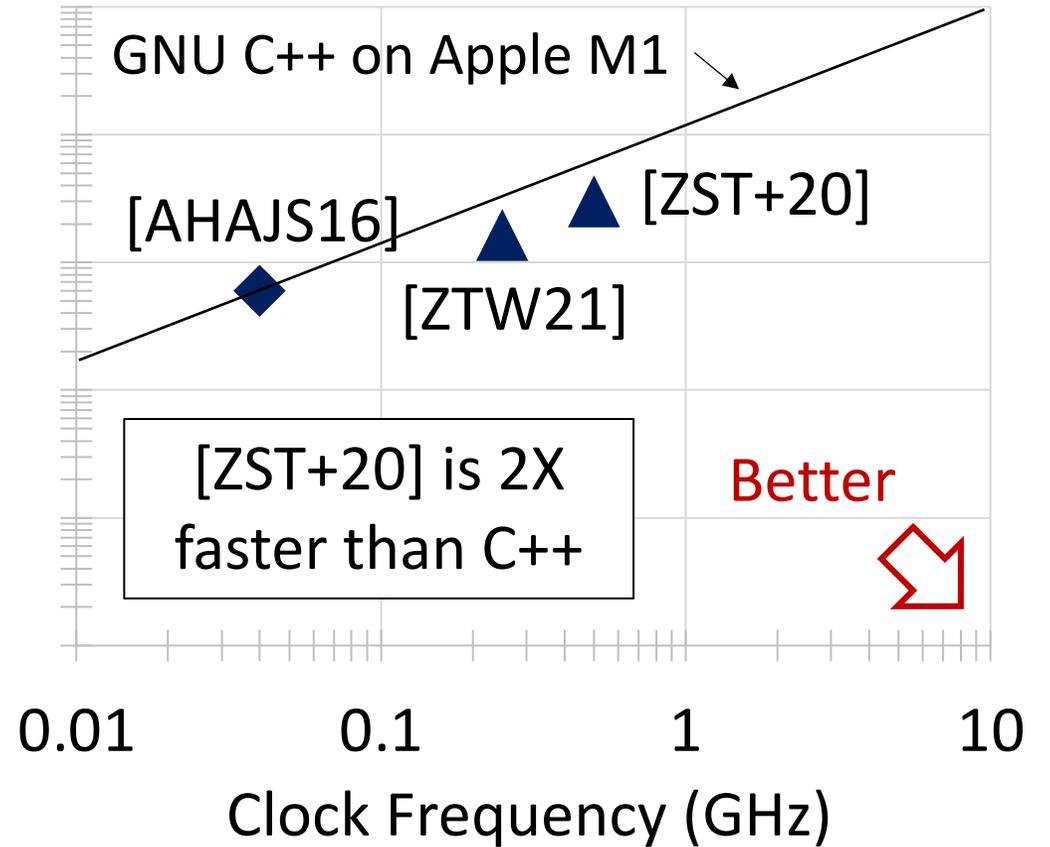
▲ ASIC



255-bit constant-time XGCD



1024-bit XGCD



● Software ◆ FPGA ▲ ASIC

Current view of XGCD design space

Target Platform

Hardware

Algorithm

÷

Application Requirements

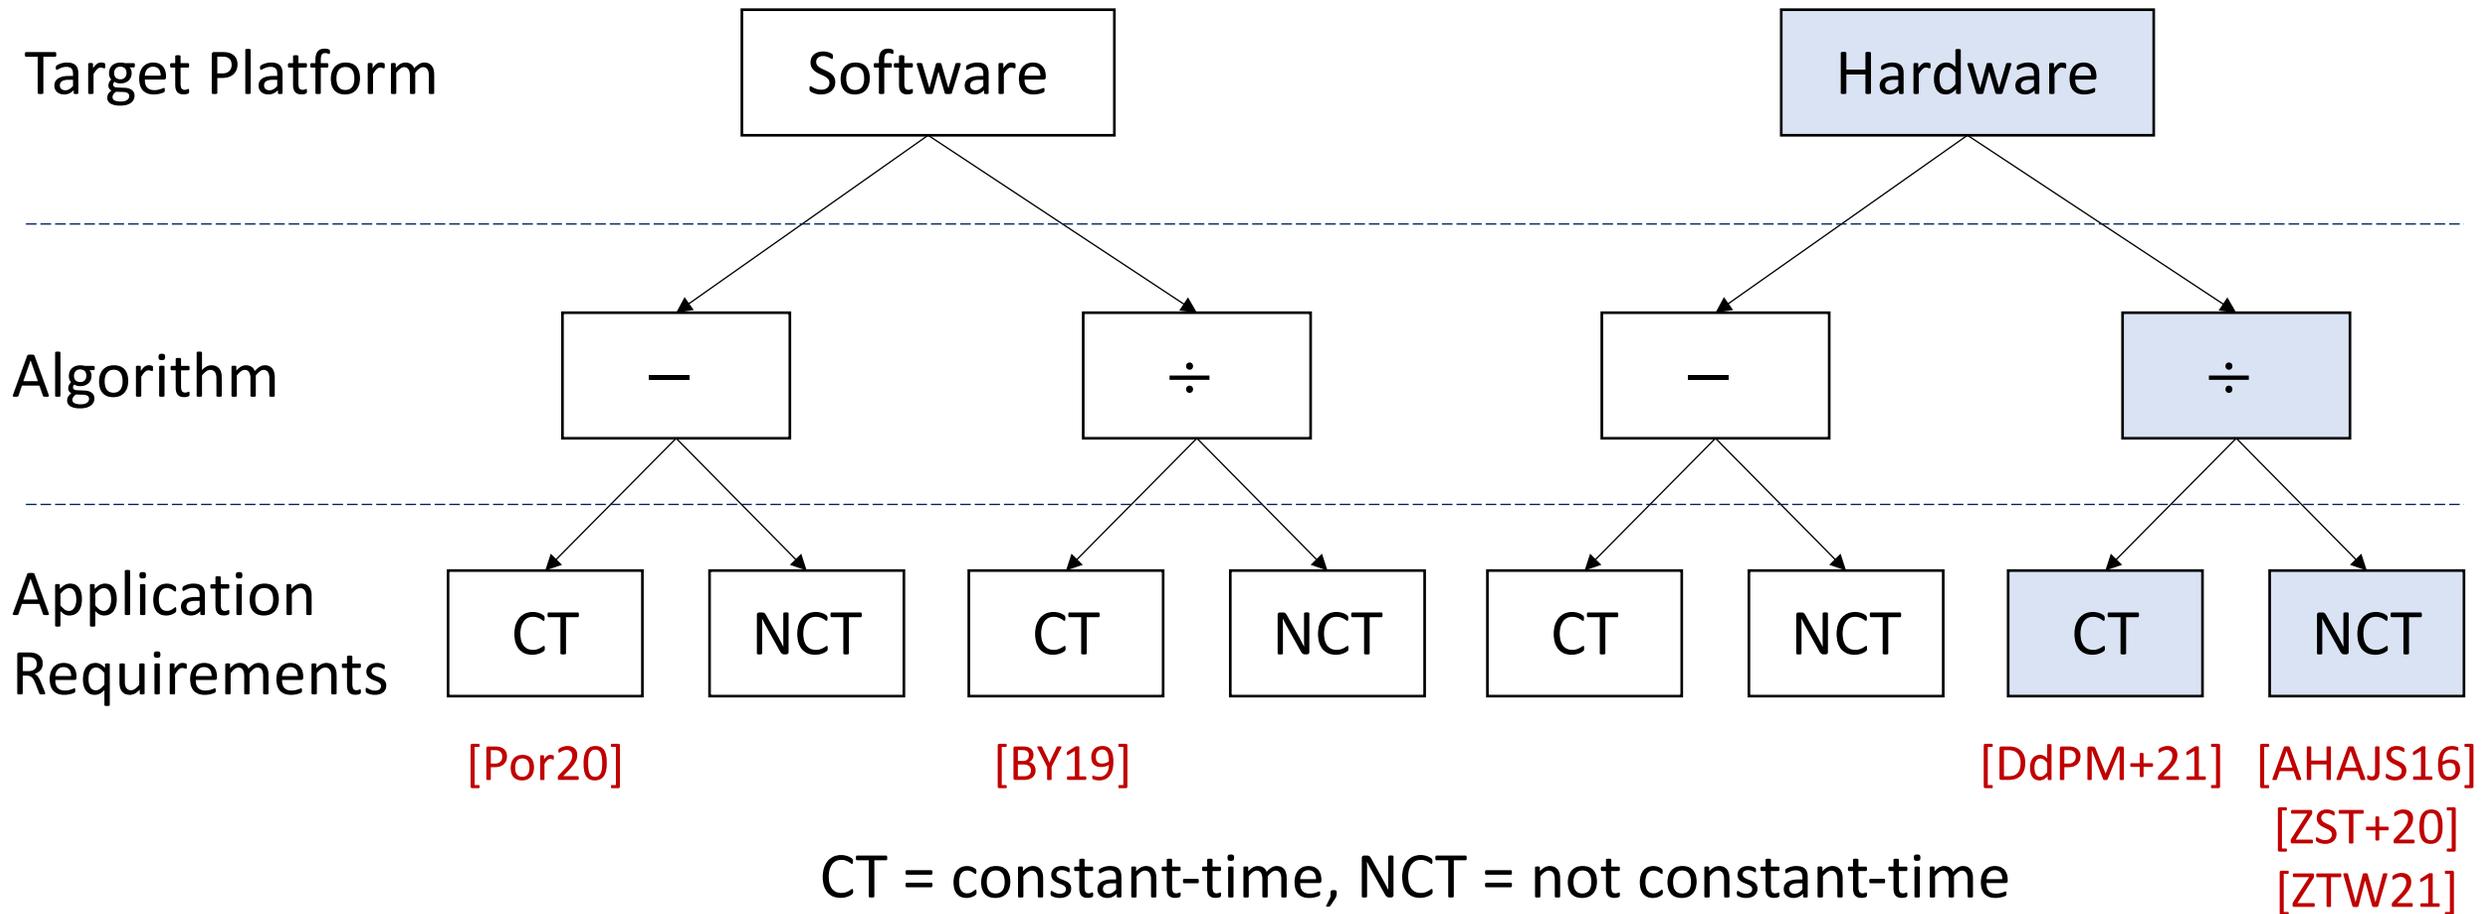
CT

NCT

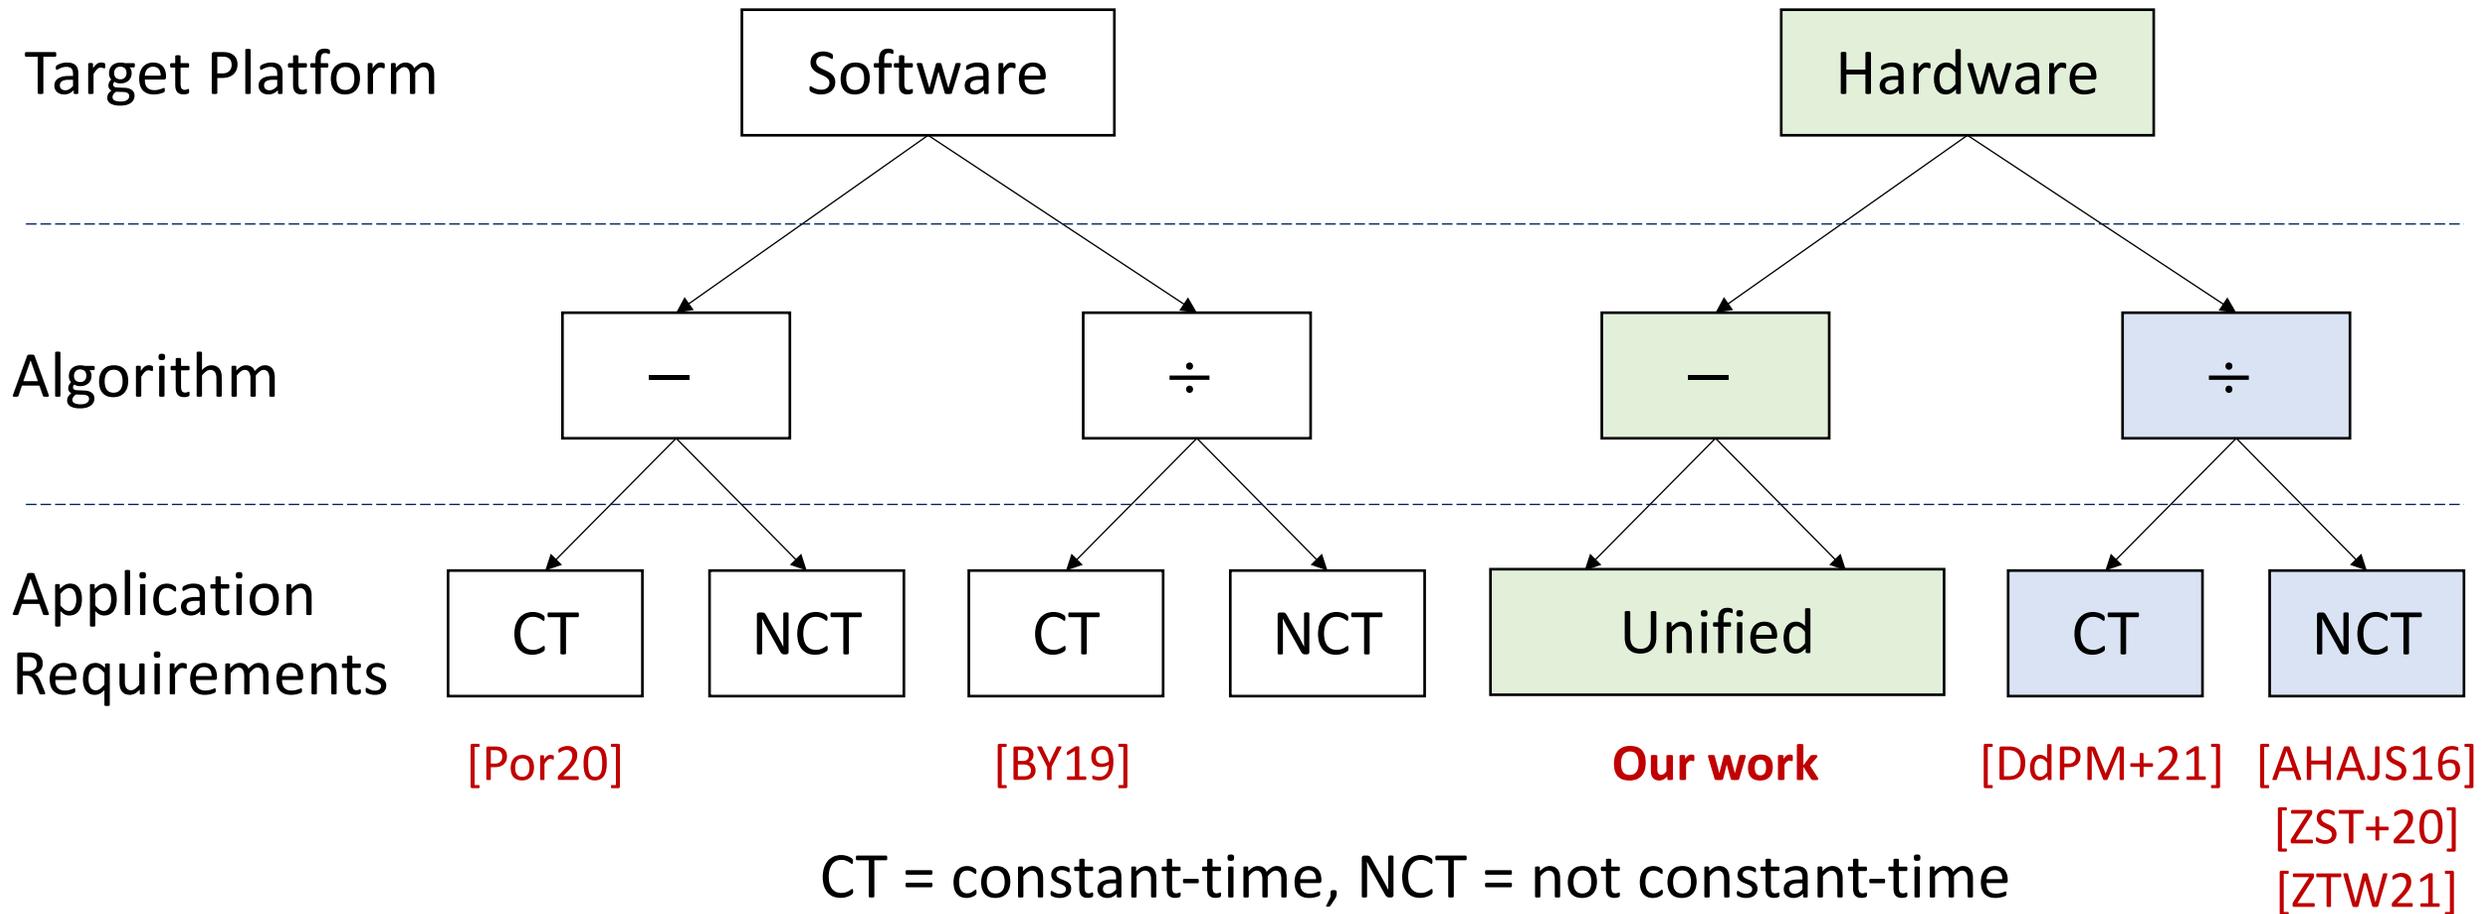
[DdPM+21] [AHAJS16]
[ZST+20]
[ZTW21]

CT = constant-time, NCT = not constant-time

We explore the broader design space



We explore the broader design space



Hardware allows for short iteration times

| Target Platform | Software | VS | Hardware |
|----------------------|----------------|----|----------------|
| Number of Iterations | From algorithm | | From algorithm |
| Constrained to ISA | Yes | | No |

Execution time = number of iterations * iteration time

The control over iteration time in hardware opens the opportunity to accelerate simpler algorithms that require more iterations.

GCD Algorithms Comparison

Stein [Ste67]

—

VS

Euclid (300 BC)

÷

GCD-preserving
Transformation

$$\gcd(a, b) = \gcd(a - b, b)$$

$$\gcd(a, b) = \gcd(a \bmod b, b)$$

GCD Algorithms Comparison

* Two-bit PM [YZ86]

Stein [Ste67]

Euclid (300 BC)

Algorithm

–

VS

÷

GCD-preserving
Transformation

$$\gcd(a, b) \\ = \gcd(a - b, b)$$

$$\gcd(a, b) \\ = \gcd(a \bmod b, b)$$

Worst-Case Iterations

387 *

1548 *

1X difference for 255 bits

1X difference for 1024 bits

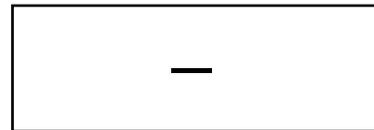
384

1542

GCD Algorithms Comparison

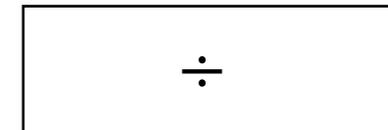
* Two-bit PM [YZ86]

Stein [Ste67]



VS

Euclid (300 BC)



| Algorithm | Stein [Ste67] | VS | Euclid (300 BC) |
|-------------------------------|-------------------------------|---|-----------------------------------|
| GCD-preserving Transformation | $\gcd(a, b) = \gcd(a - b, b)$ | | $\gcd(a, b) = \gcd(a \bmod b, b)$ |
| Worst-Case Iterations | 387 * 1548 * | 1X difference for 255 bits 1X difference for 1024 bits | 384 1542 |
| Average Iterations | 300 * 1195 * | 1.6X difference for 255 bits 2X difference for 1024 bits | 189 598 |

From GCD to Extended GCD (XGCD)

- Compute Bézout coefficients satisfying Bézout Identity

$$\mathbf{b}_a, \mathbf{b}_b : \mathbf{b}_a * a_0 + \mathbf{b}_b * b_0 = \text{gcd}(a_0, b_0)$$

- Maintain these relations each cycle, where $\text{gcd}(a_0, b_0) = \text{gcd}(a, b)$

$$\begin{aligned} u * a_0 + m * b_0 &= a \\ y * a_0 + n * b_0 &= b \end{aligned}$$

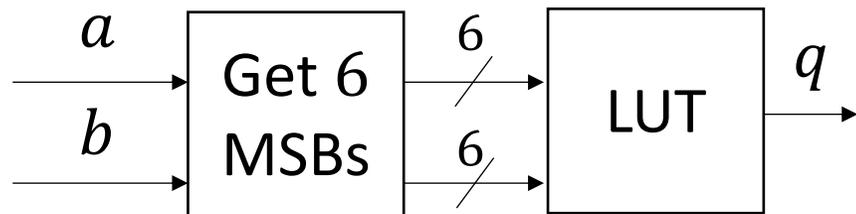
Two-bit PM Critical Path

GCD update: $a = \frac{a-b}{4}$

XGCD update: $m = \frac{m-n-a_m}{4}$

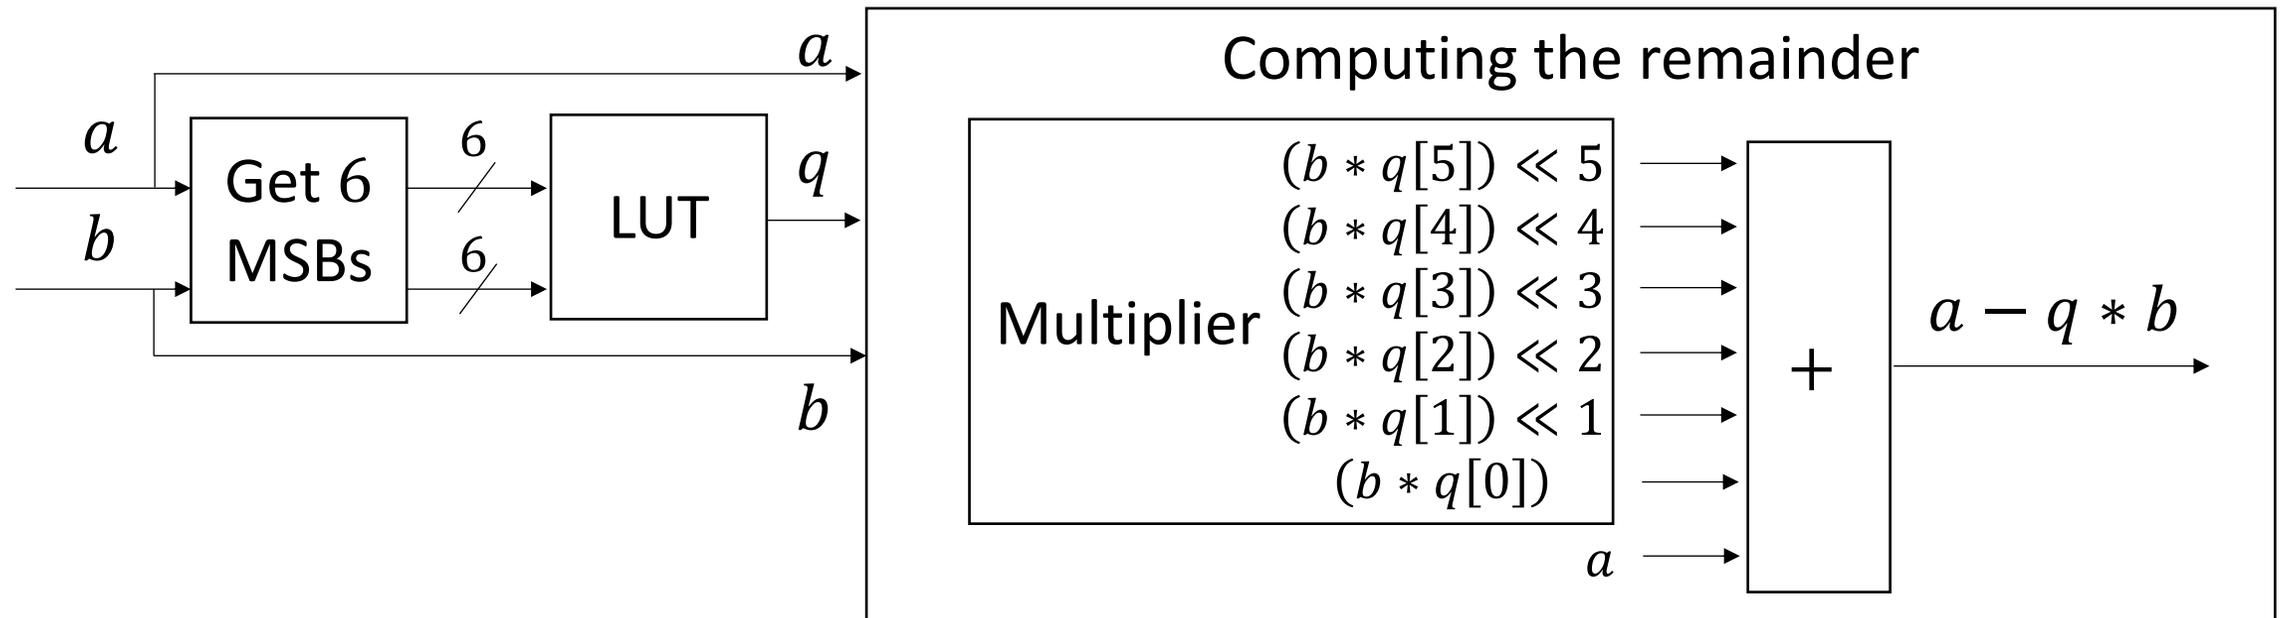
Euclid critical path

Compute $q \leq \lfloor \frac{a}{b} \rfloor$ \longrightarrow Compute $q * b$ \longrightarrow Compute $a - q * b$



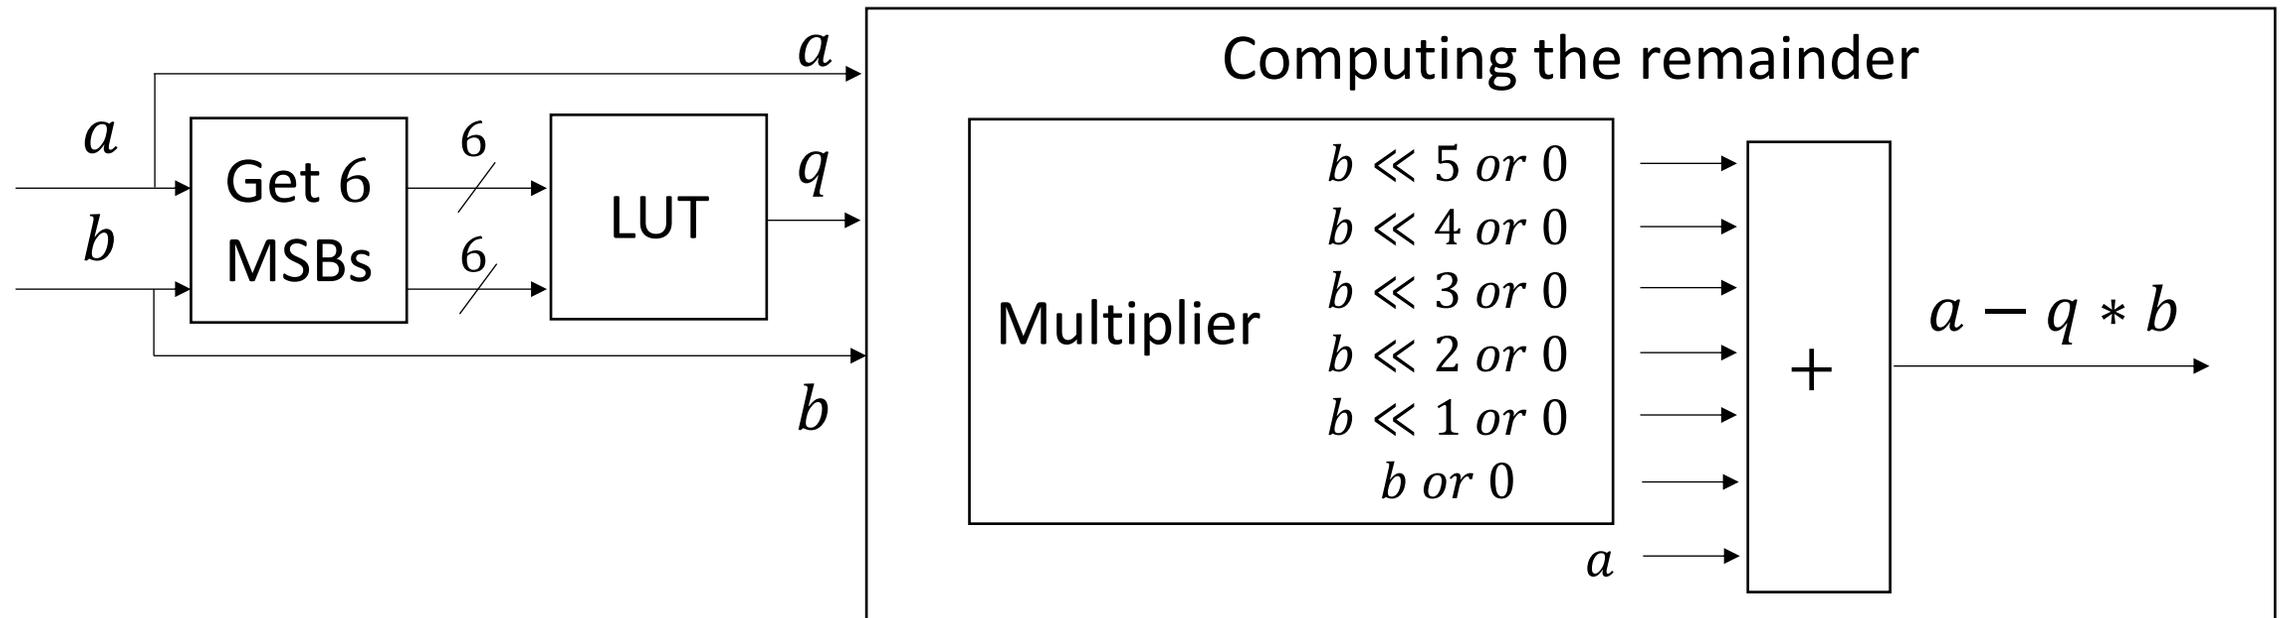
Euclid critical path

Compute $q \leq \lfloor \frac{a}{b} \rfloor$ \longrightarrow Compute $q * b$ \longrightarrow Compute $a - q * b$



Euclid critical path

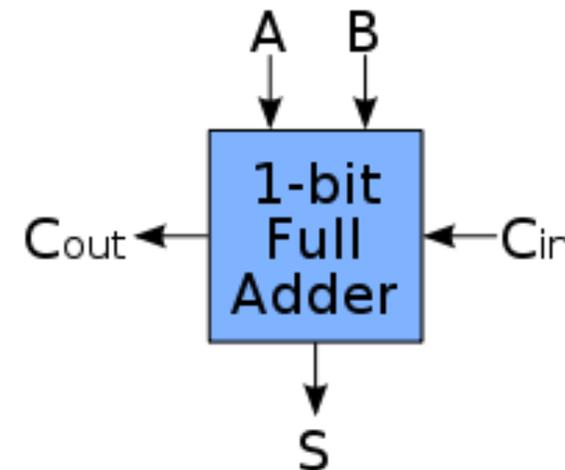
Compute $q \leq \lfloor \frac{a}{b} \rfloor$ \longrightarrow Compute $q * b$ \longrightarrow Compute $a - q * b$



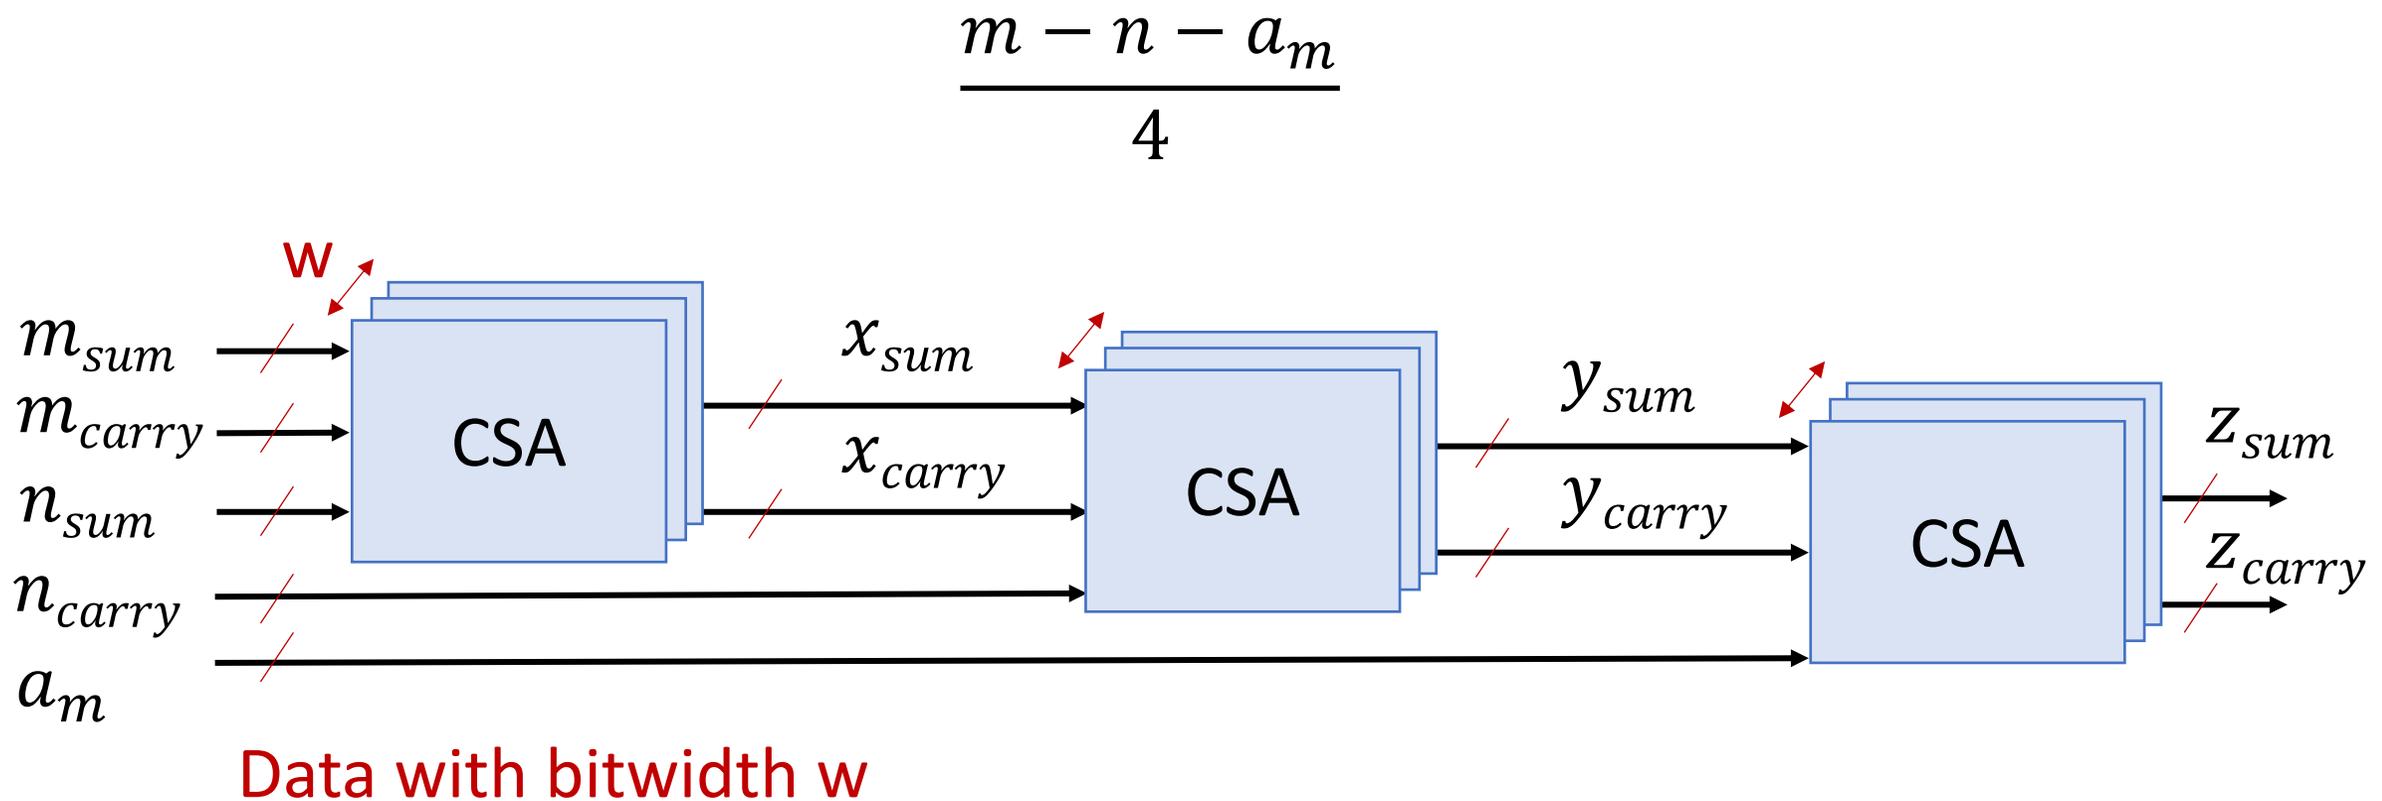
Critical paths primarily require additions

- The fastest adder is a carry-save adder (CSA)
 - Eliminates carry propagation, requiring $O(1)$ delay
 - Stores numbers in CSA form or redundant binary form

$$\begin{array}{r}
 \color{red}{1110} \\
 1101 \text{ (a)} \\
 + 1110 \text{ (b)} \\
 0010 \text{ (c)} \\
 \hline
 11101
 \end{array}
 \longrightarrow
 \begin{array}{r}
 1101 \text{ (a)} \\
 + 1110 \text{ (b)} \\
 0010 \text{ (c)} \\
 \hline
 + 0001 \text{ sum} \\
 \color{red}{11100} \text{ carry} \\
 \hline
 11101
 \end{array}$$

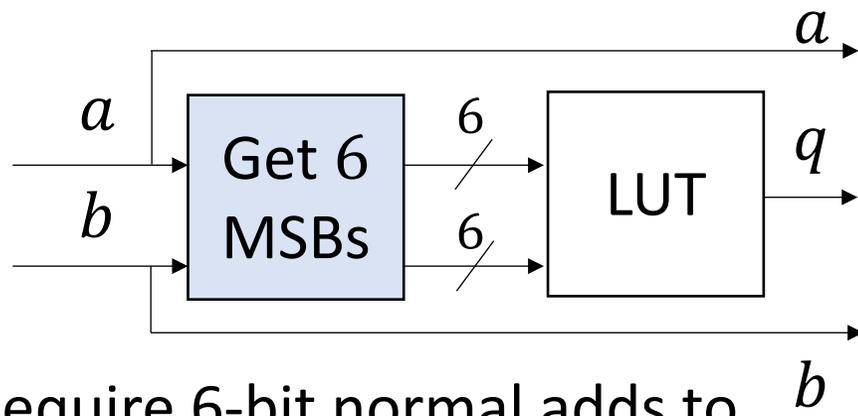


Two-bit PM critical path: 3 CSA delays



Euclid critical path

Compute $q \leq \lfloor \frac{a}{b} \rfloor$ \longrightarrow Compute $q * b$ \longrightarrow Compute $a - q * b$

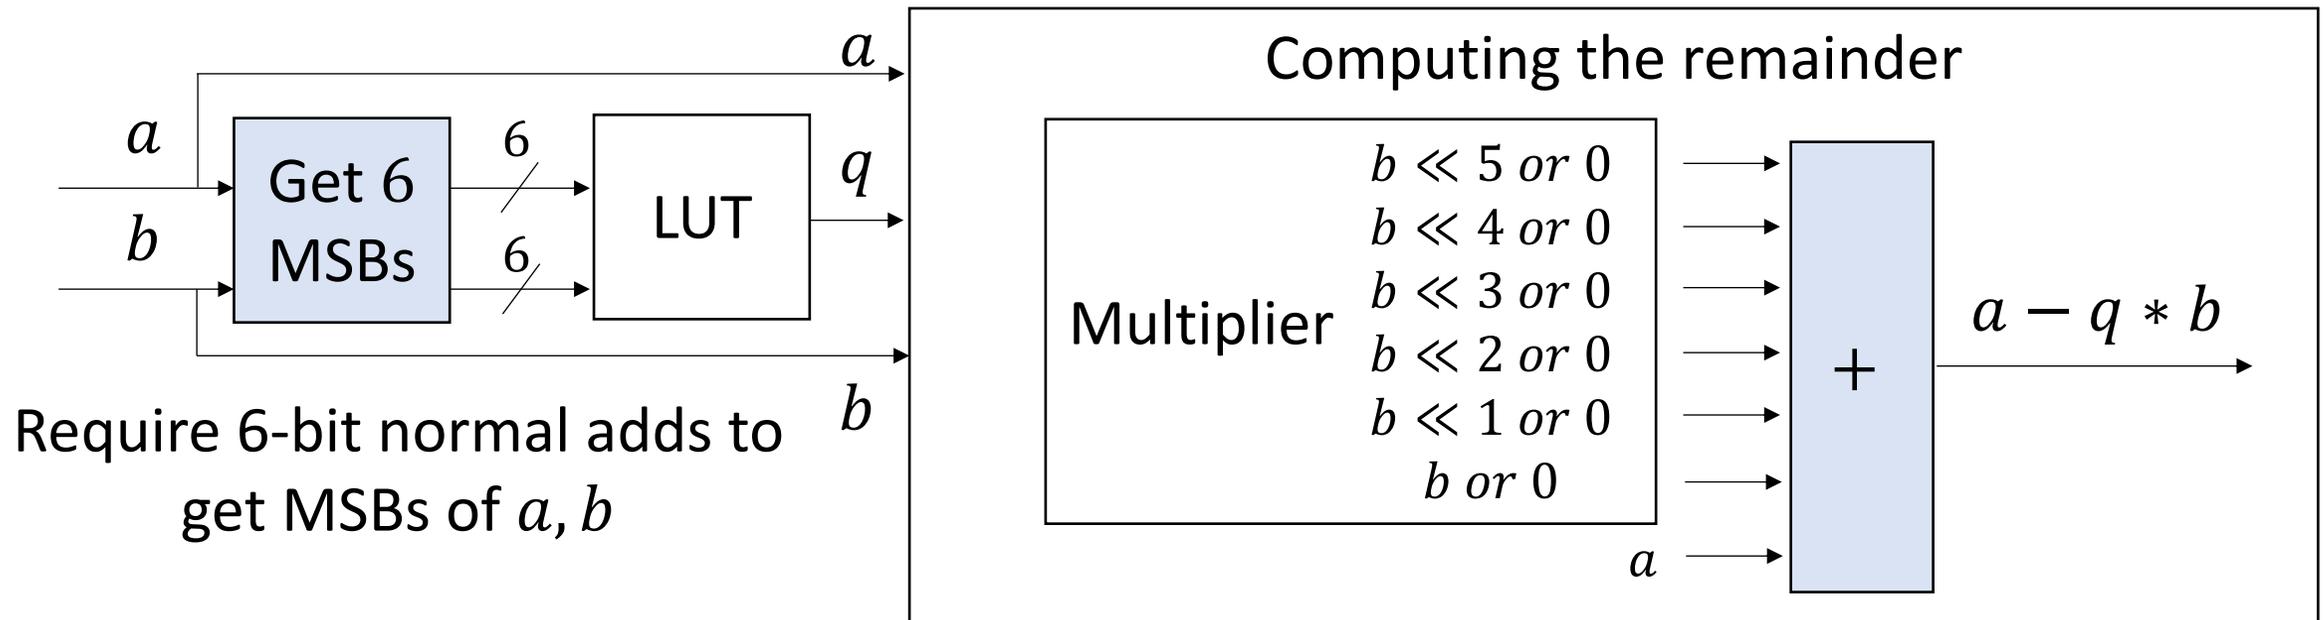


Require 6-bit normal adds to
get MSBs of a, b

$\lfloor \log_2(6) \rfloor + 1 = 3$ CSA delays

Euclid critical path: at least 9 CSA delays

Compute $q \leq \lfloor \frac{a}{b} \rfloor$ \longrightarrow Compute $q * b$ \longrightarrow Compute $a - q * b$



$\lfloor \log_2(6) \rfloor + 1 = 3$ CSA delays

Add 14 values with CSAs $\approx \lfloor \log_{3/2}(14) \rfloor = 6$ CSA delays

Two-bit PM is a faster starting point

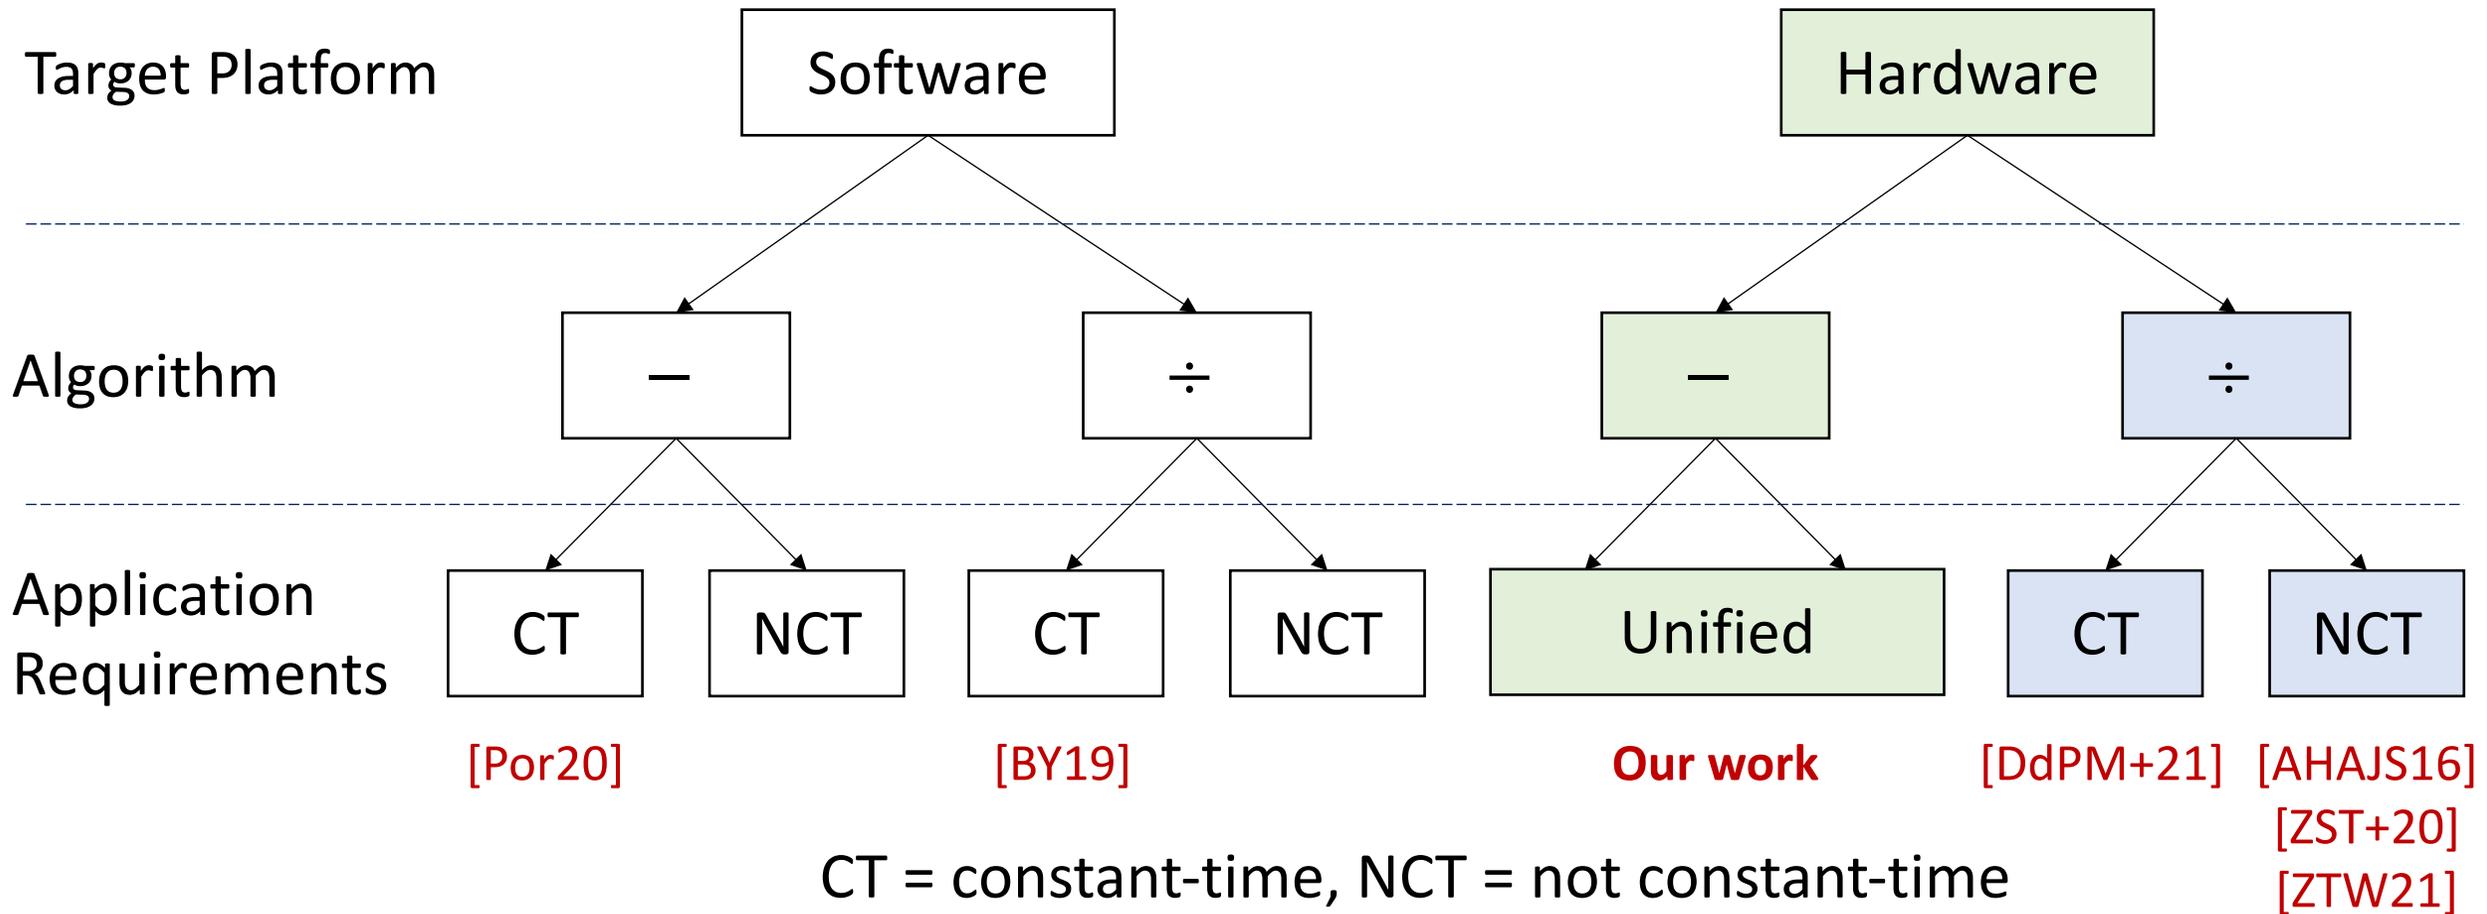
- Two-bit PM critical path is at least 3X shorter than Euclid's
- Two-bit PM iteration counts are at most 2X higher than Euclid's

Two-bit PM with carry-save adders is the more promising starting point for hardware in the average and the worst-case.

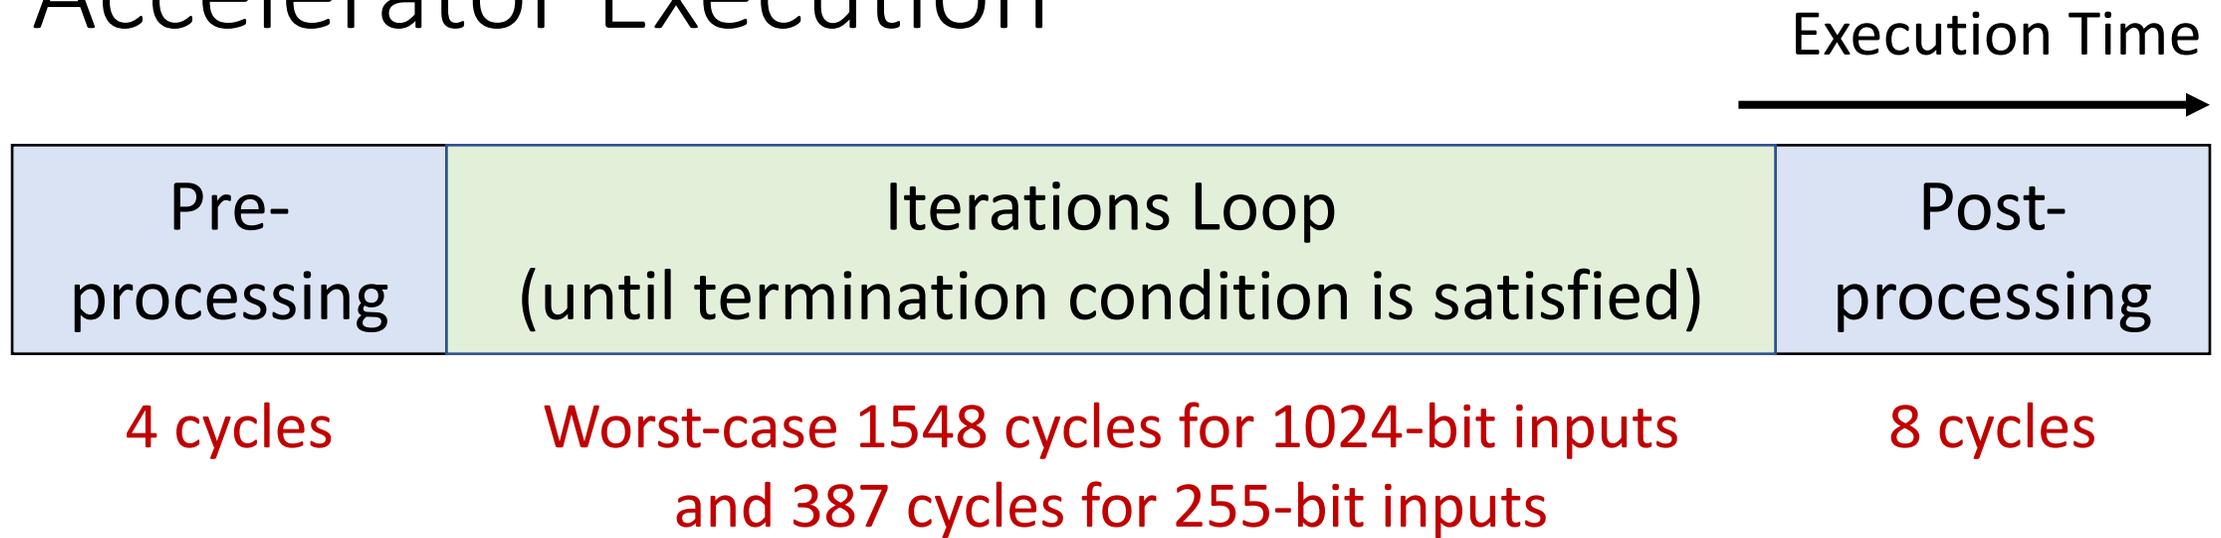
Our unified design with constant-time config

| Application Requirements | CT | VS | NCT |
|---|---------------------------------|----|-------------------------|
| Approach | Pad to worst-case cycle count | | Reduce inputs until GCD |
| Termination Condition | Cycle count equal to worst case | | $a == 0$ or $b == 0$ |
| <p>Note that since a, b are in CSA form, we do not know when they become 0</p> | | | |

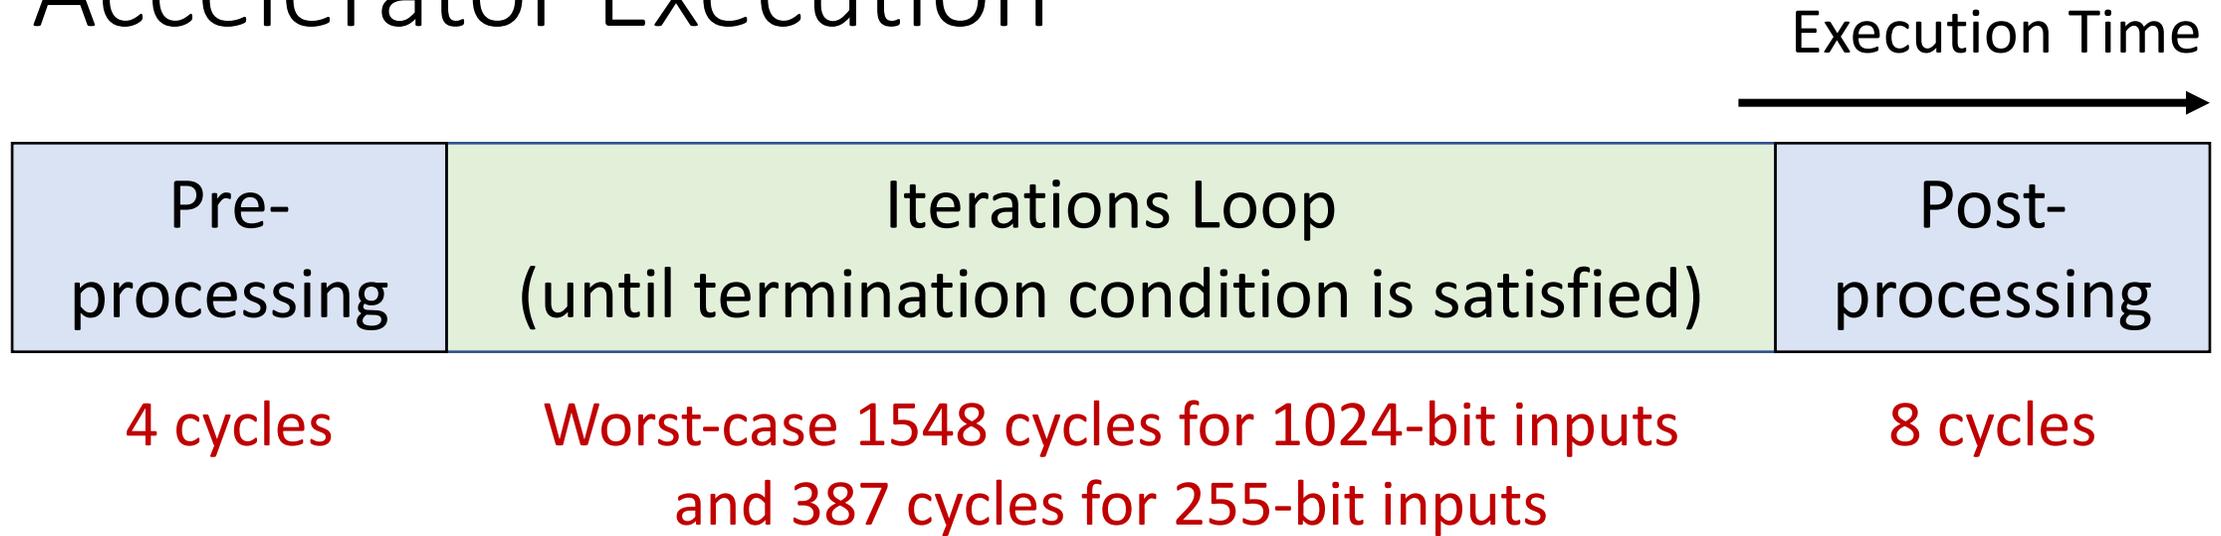
We focus on the optimal design space



Accelerator Execution

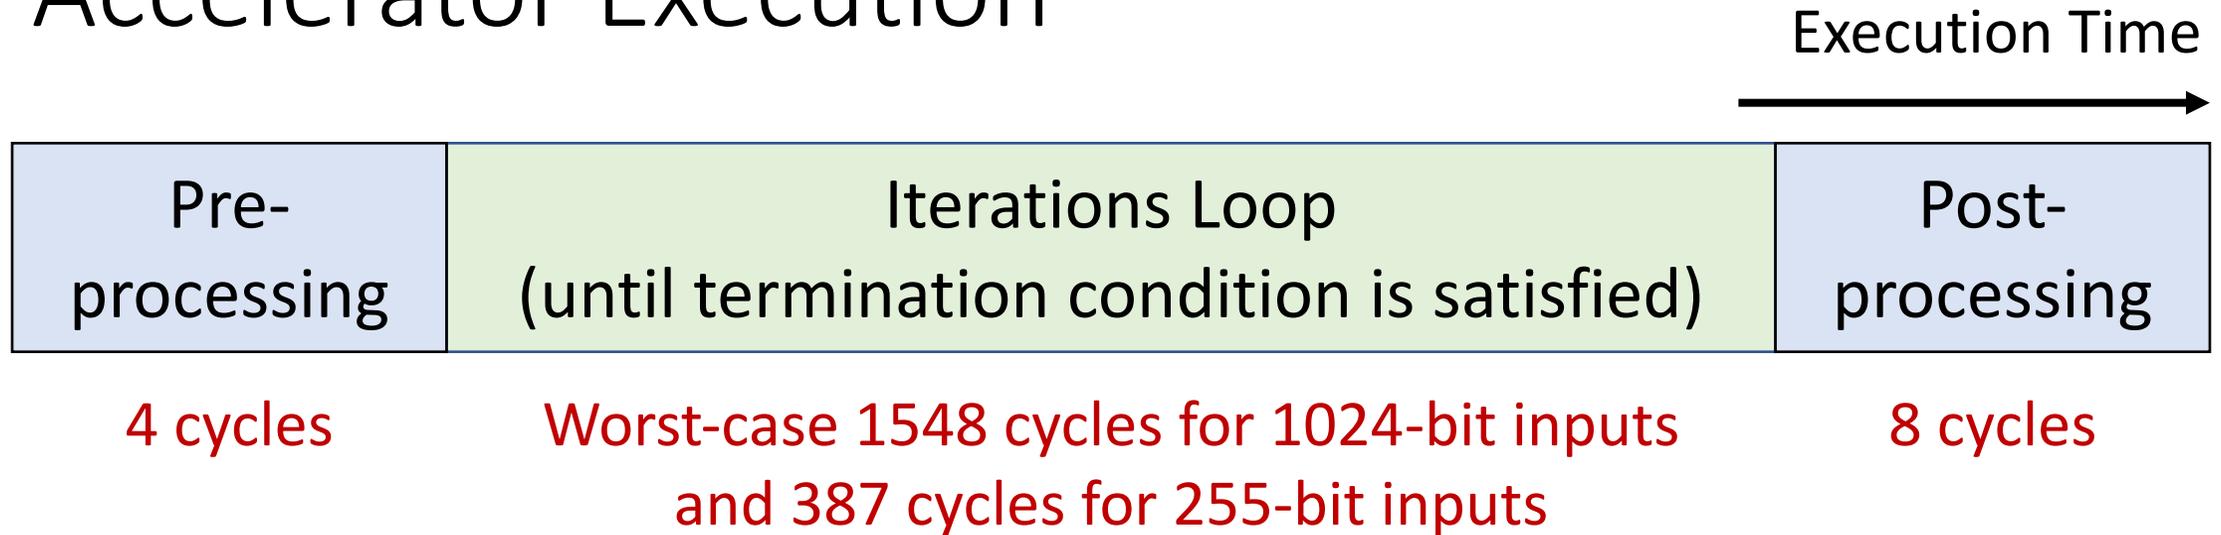


Accelerator Execution



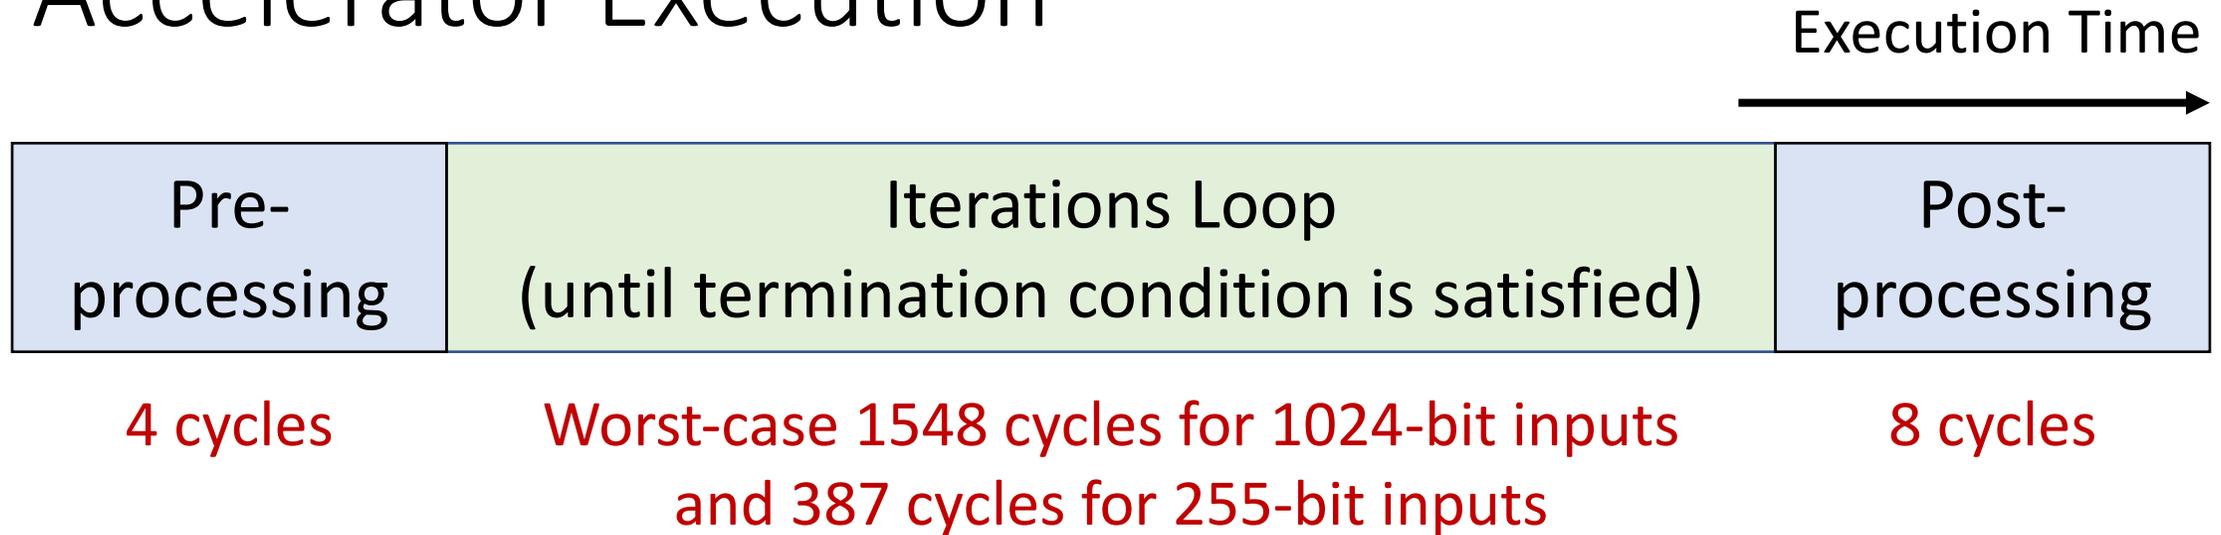
- Preserve results when shifting in CSA form

Accelerator Execution



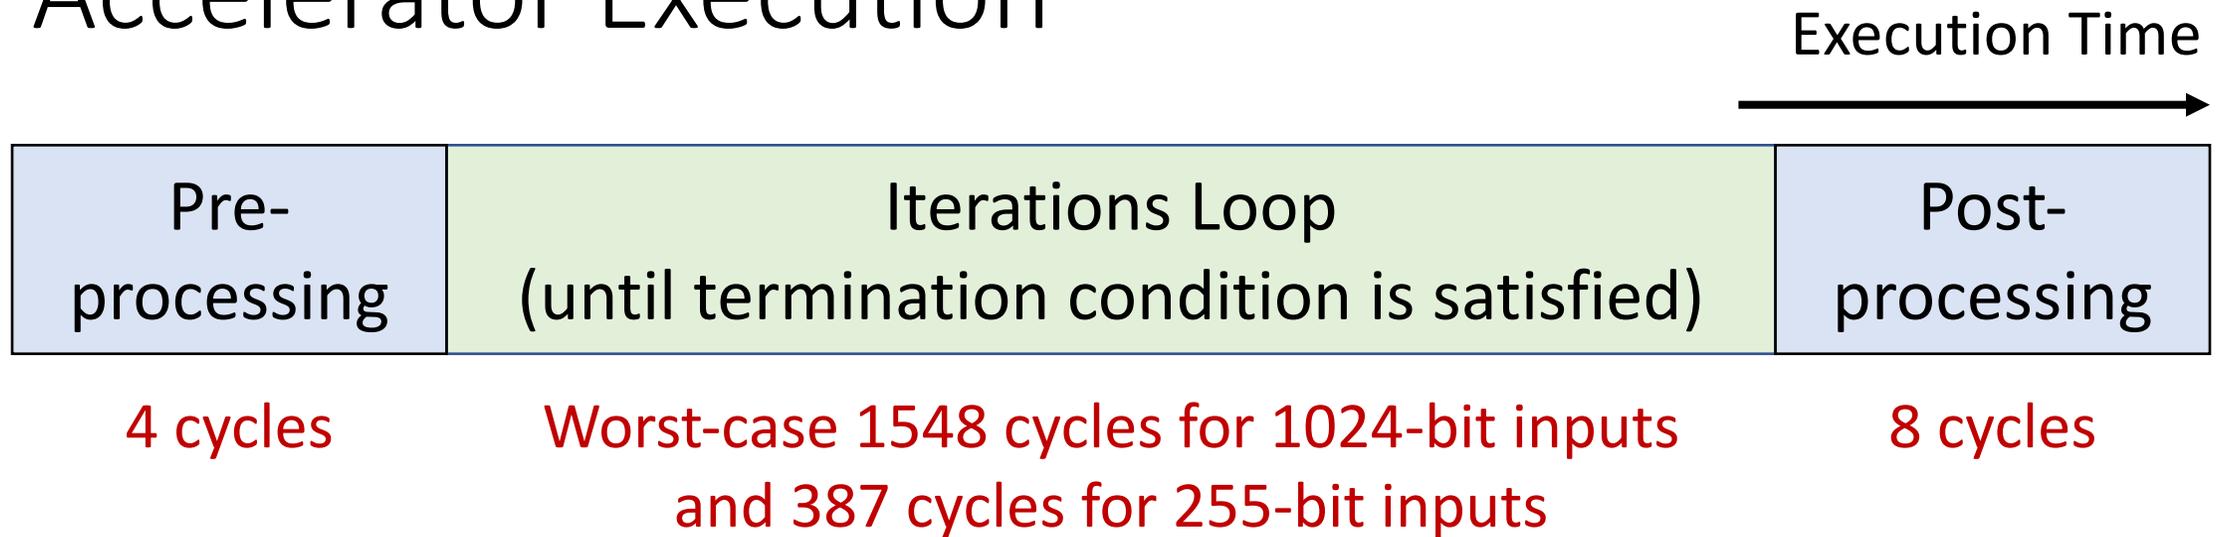
- Preserve results when shifting in CSA form
- Allocate multiple cycles for processing steps

Accelerator Execution



- Preserve results when shifting in CSA form
- Allocate multiple cycles for processing steps
- Subsample a, b for termination condition

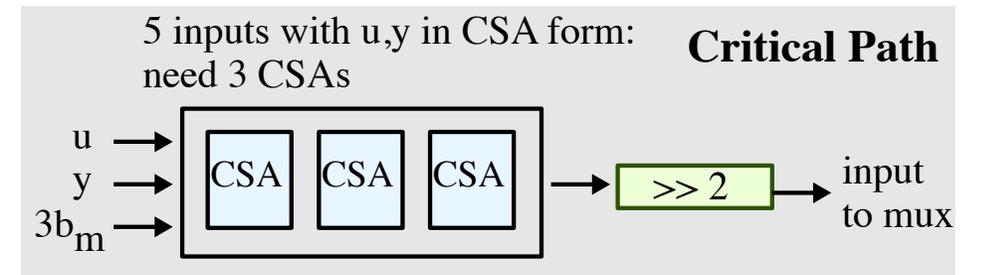
Accelerator Execution



- Preserve results when shifting in CSA form
- Allocate multiple cycles for processing steps
- Subsample a, b for termination condition
- Minimize control overhead

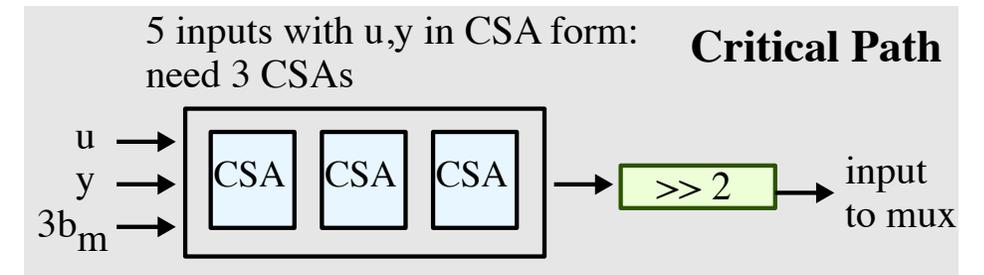
Critical Path for ASIC in 16nm

| | 255-bit XGCD | 1024-bit XGCD |
|----------------------|--------------|---------------|
| DFF clk to Q | 45 | 40 |
| Inverter | 7 | 0 |
| CSA | 18 | 39 |
| CSA | 31 | 39 |
| Buffer | 13 | 0 |
| CSA | 30 | 34 |
| Shift in CSA form | 15 | 18 |
| Late select muxes | 18 | 18 |
| Precomputing control | 27 | 22 |
| Setup Time | 2 | 5 |
| Total | 204 | 215 |



Critical Path for ASIC in 16nm

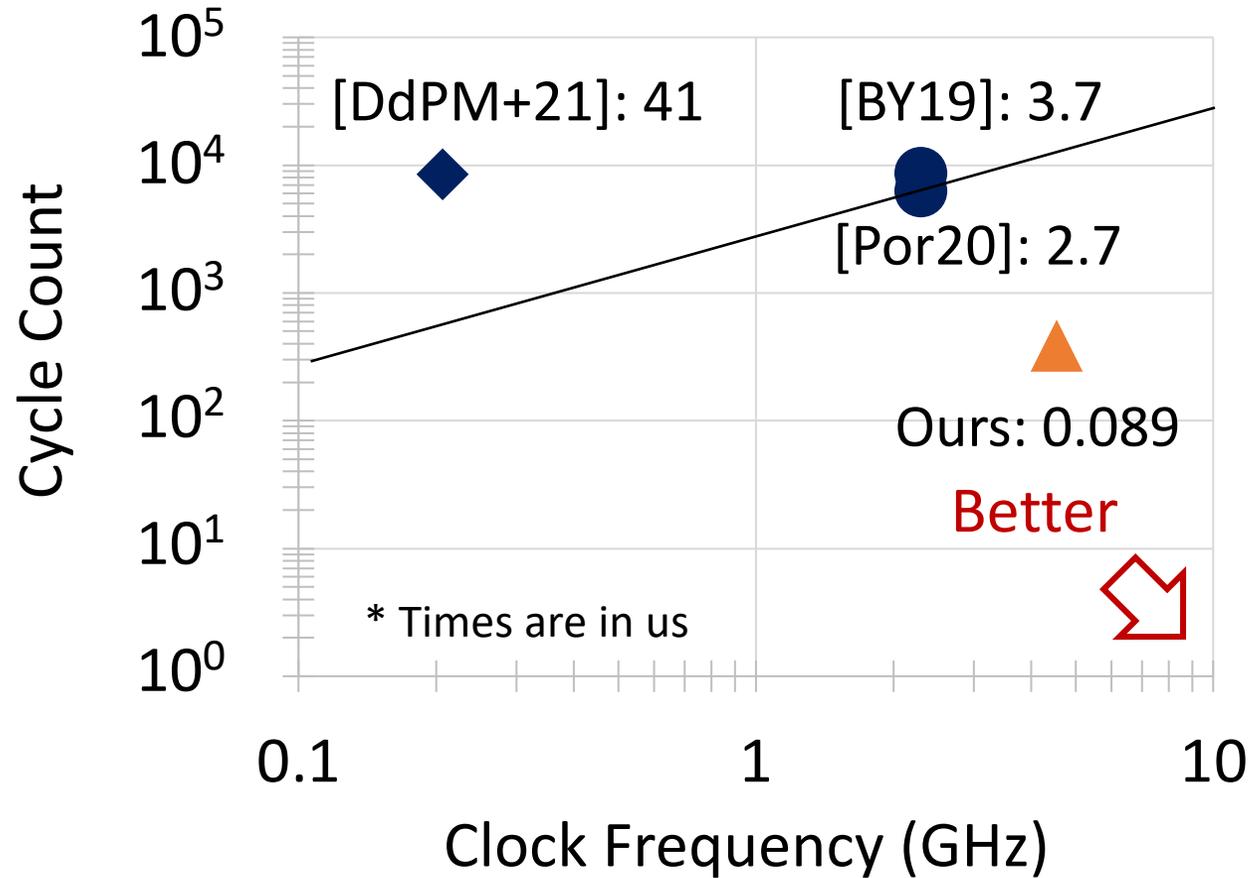
| | 255-bit XGCD | 1024-bit XGCD |
|----------------------|--------------|---------------|
| DFF clk to Q | 45 | 40 |
| Inverter | 7 | 0 |
| CSA | 18 | 39 |
| CSA | 31 | 39 |
| Buffer | 13 | 0 |
| CSA | 30 | 34 |
| Shift in CSA form | 15 | 18 |
| Late select muxes | 18 | 18 |
| Precomputing control | 27 | 22 |
| Setup Time | 2 | 5 |
| Clock Skew | 16 | 41 |
| Total | 220 | 257 |



These are post-layout numbers for a fabrication-ready design

255-bit XGCD: 4.5 GHz
1024-bit XGCD: 3.9 GHz

255-bit Constant-time XGCD Comparison



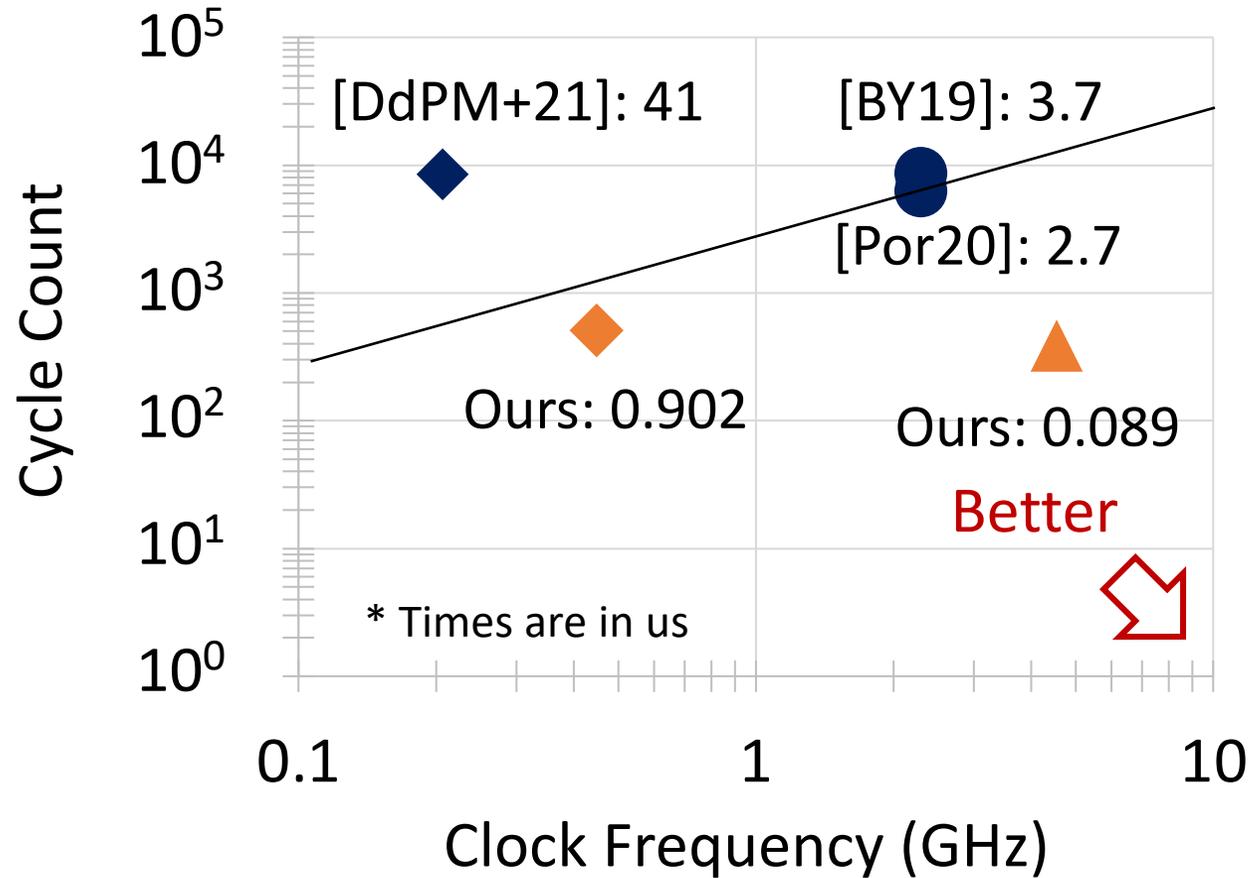
- Our ASIC
- 31X faster than [Por20]
 - First for constant-time 255-bit XGCD

● Software

◆ FPGA

▲ ASIC

255-bit Constant-time XGCD Comparison

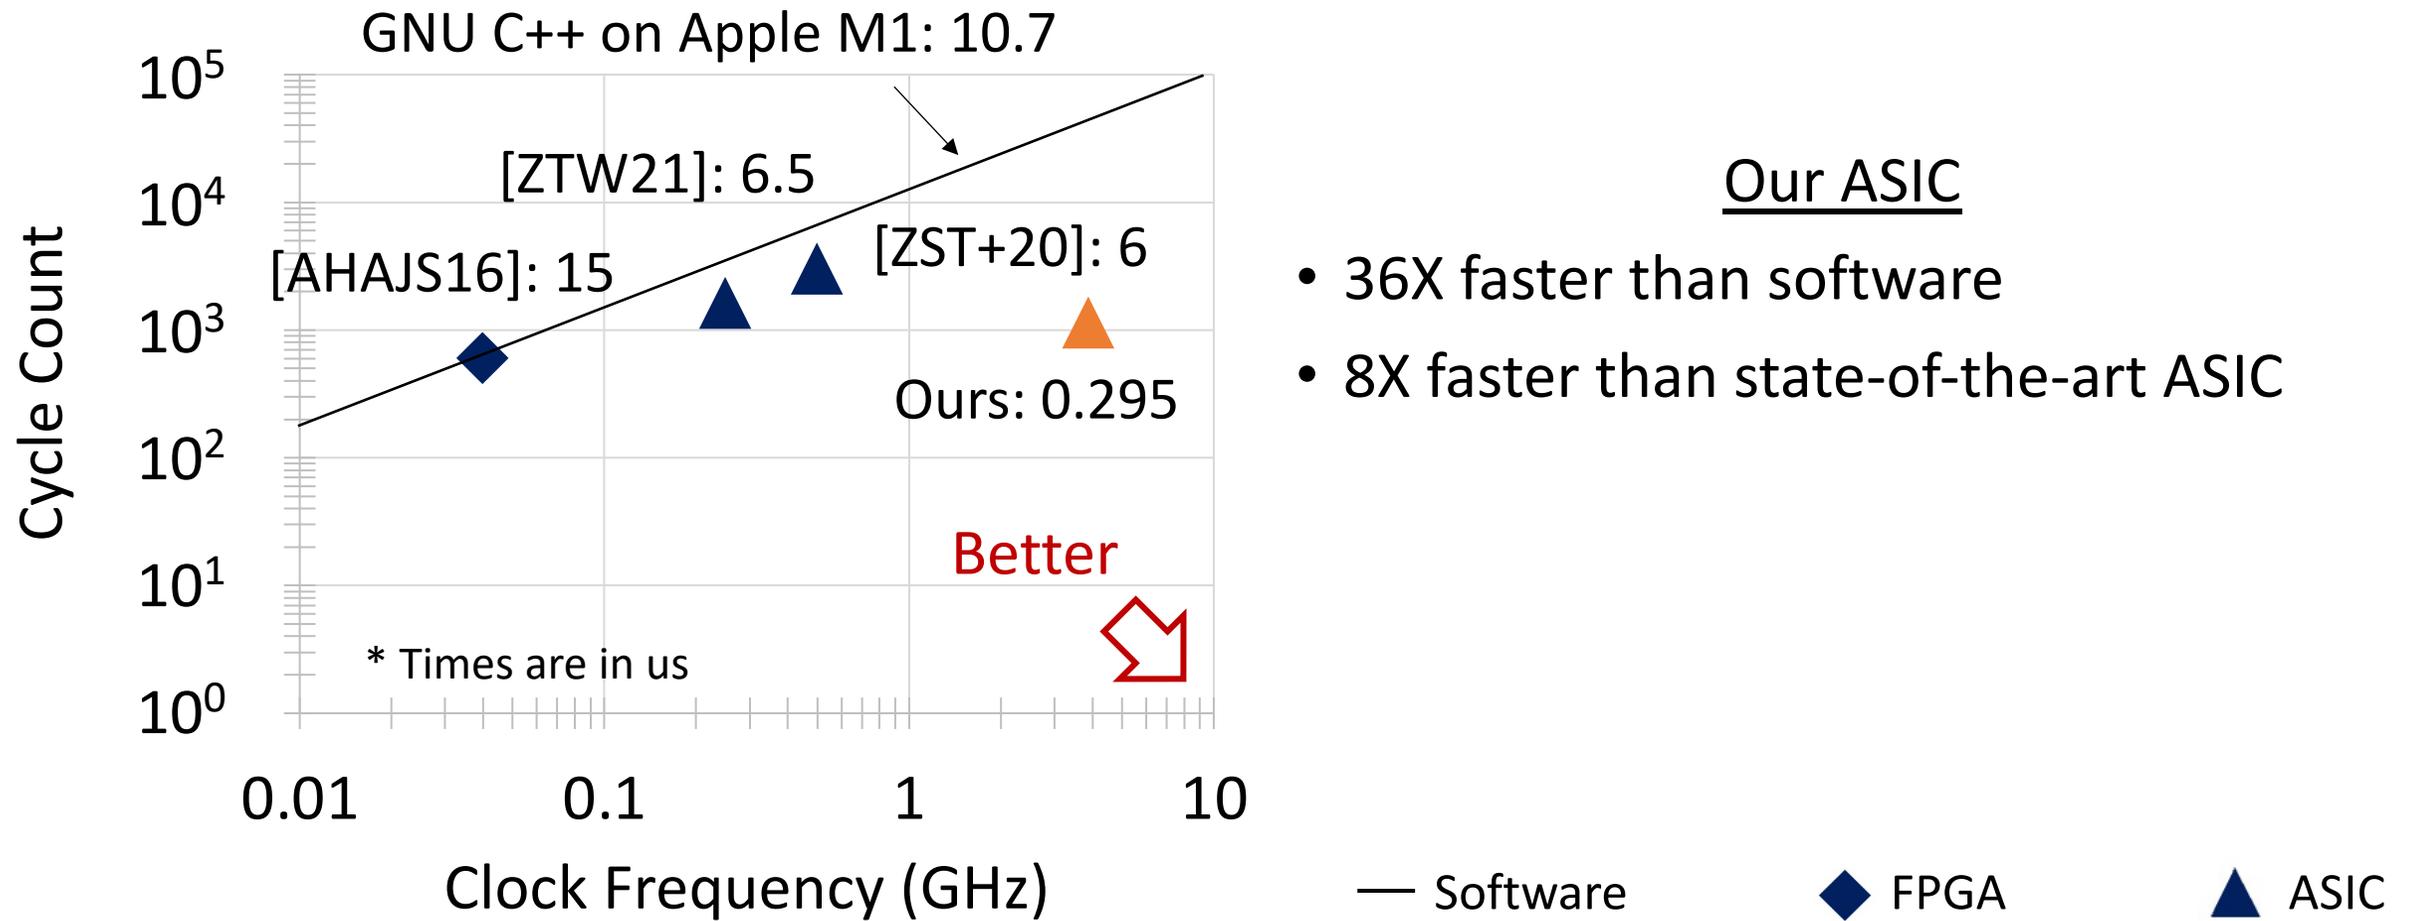


- Our ASIC
- 31X faster than [Por20]
 - First for constant-time 255-bit XGCD

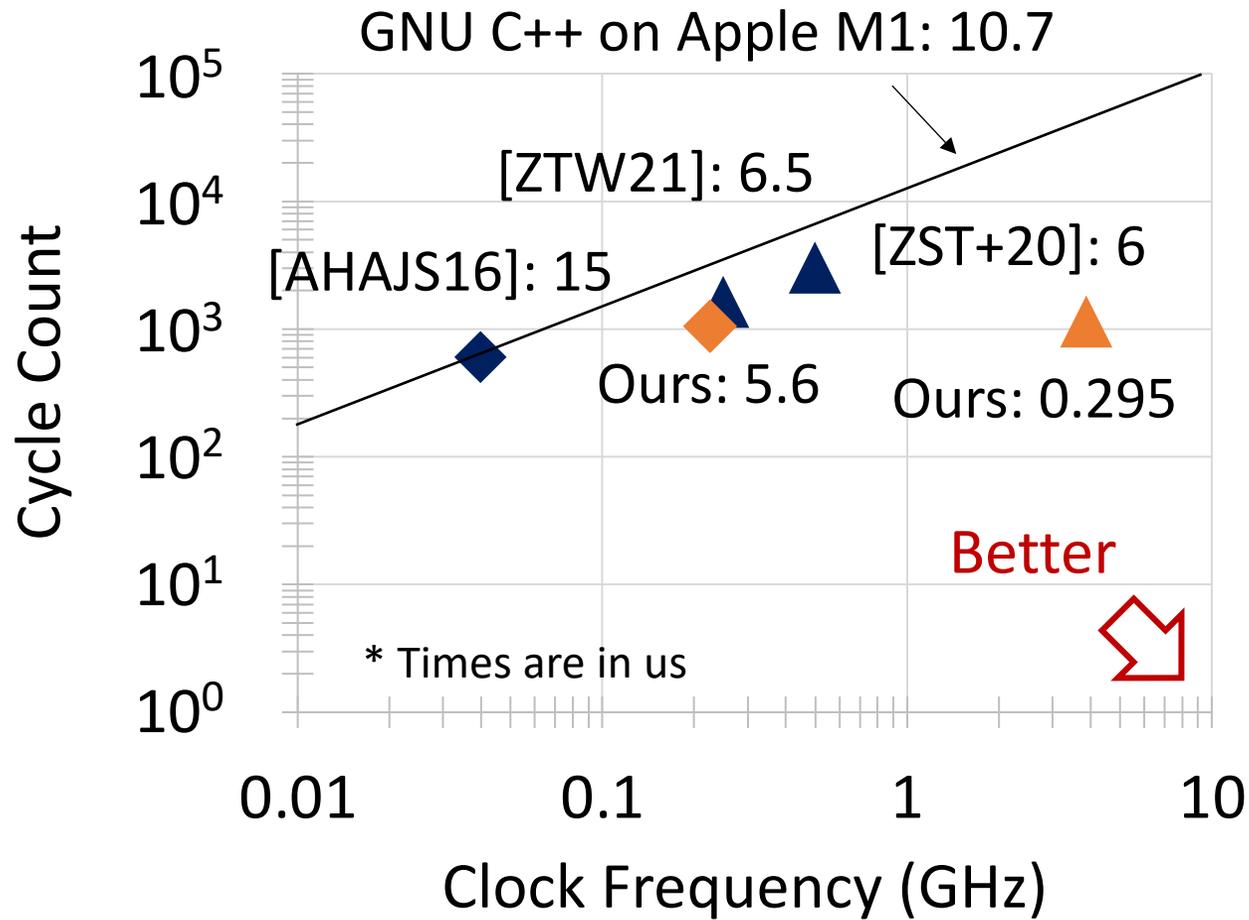
Direct FPGA Comparison

Our design is 45X faster

1024-bit XGCD Comparison



1024-bit XGCD Comparison



- Our ASIC
- 36X faster than software
 - 8X faster than state-of-the-art ASIC

Direct FPGA Comparison

Our design is 2.7X faster

— Software

◆ FPGA

▲ ASIC

Our design impacts application approaches

1. Supports progression in state of the art for Curve25519

Our design impacts application approaches

1. Supports progression in state of the art for Curve25519
2. Informs reasonable security levels for this type of VDF [ZZOTW23]

Our design impacts application approaches

1. Supports progression in state of the art for Curve25519
2. Informs reasonable security levels for this type of VDF [ZZOTW23]
3. May be useful for other applications

Our design impacts application approaches

1. Supports progression in state of the art for Curve25519
2. Informs reasonable security levels for this type of VDF [ZZOTW23]
3. May be useful for other applications



<https://github.com/kavyasreedhar/sreedhar-xgcd-hardware-ches2022>