Toward Fair and Comprehensive Benchmarking of CAESAR Candidates in Hardware: Standard API, High-Speed Implementations in VHDL/Verilog, and Benchmarking Using FPGAs

Ekawat Homsirikamol,
William Diehl, Ahmed Ferozpuri,
Farnoud Farahmand,
Michael X. Lyons, Panasayya Yalla,
and Kris Gaj
George Mason University
USA

Based on work partially supported by the
National Science Foundation
under Grant No. 1314540

GMU Benchmarking Team

"Ice" Homsirikamol

Will Diehl

Ahmed Ferozpuri

Farnoud Farahmand

Mike X. Lyons

Panasayya Yalla

Evaluation Criteria in Cryptographic Contests

Security

Software Efficiency

µProcessors µControllers

Hardware Efficiency

FPGAs ASICs

Flexibility

Simplicity

Licensing

Hardware Benchmarking in Previous Contests

AES (1999-2000): 5 final candidates

eSTREAM (2007-2008): 8 Phase-3 candidates

SHA-3 (2010-2012): 14 Round 2 Candidates
+ 5 Final Candidates

CAESAR (2016): 29 Round 2 Candidates

New in CAESAR

1) standard hardware **Application Programming Interface (API)**
2) comprehensive **Implementer's Guide** and **Development Package**, including VHDL and Python code common for all candidates
3) the design teams have been asked to submit their **own Verilog/VHDL code**

CAESAR
Hardware API

CAESAR Hardware API

Specifies:

- Minimum Compliance Criteria
- Interface
- Communication Protocol
- Timing Characteristics

Assures:

- Compatibility
- Fairness

CAESAR Hardware API - Timeline

- July 2015, CryptArchi, Leuven, GMU API v1.0
- Sep. 2015, DIAC, Singapore, GMU API v1.1
- Dec. 2015, ReConFig, Cancun, GMU API v1.2

- Feb. 16, 2016, proposed CAESAR API v1.0
- Mar. 22, 2016, CAESAR Committee considers adoption
- May 7, 2016, official adoption by the CAESAR Committee
- May 12, 2016, final version of CAESAR API v1.0

- June 30, 2016, deadline for VHDL/Verilog Code
- August 12, 2016, last submission of the code

CAESAR API v1.0 vs. GMU API v1.2

Feb. 16, 2016

- Functional Changes
  - Supporting both high-speed and lightweight implementations
  - Supporting both single-pass and two-pass algorithms
  - **Moving the buffering of decrypted data to an external unit, common for all candidates**
  - **No passing of Npub and AD to the output**
  - **Specifying the maximum size of AD/message/ciphertext explicitly**
  - **Requiring full support for key scheduling**

- Editorial Changes
  - Adding **Minimum Compliance Criteria** & Timing Characteristics
  - Separating from the Implementer's Guide

Advantages of CAESAR API v1.0 vs. GMU API 1.2

- **Simplified:**
  - code development
  - definitions of timing parameters for decryption
  - resource utilization characterization
  - benchmarking
- **Aimed to**
  - speed-up coding
  - encourage more design teams to get involved

Limitations of the CAESAR API v1.0

Interface:

- **No parallel loading of AD and Message**
  (used by Keyak)

Protocol:

- **No support for intermediate tags**
  (used by variants of ELmD, POET, TriviA-ck, and COLM)

- **No protocol support for a second pass without storing intermediate results (or the entire input) inside of the authenticated cipher core**

CAESAR
Implementer's Guide & Development Package

Top-level block diagram of a High-Speed architecture

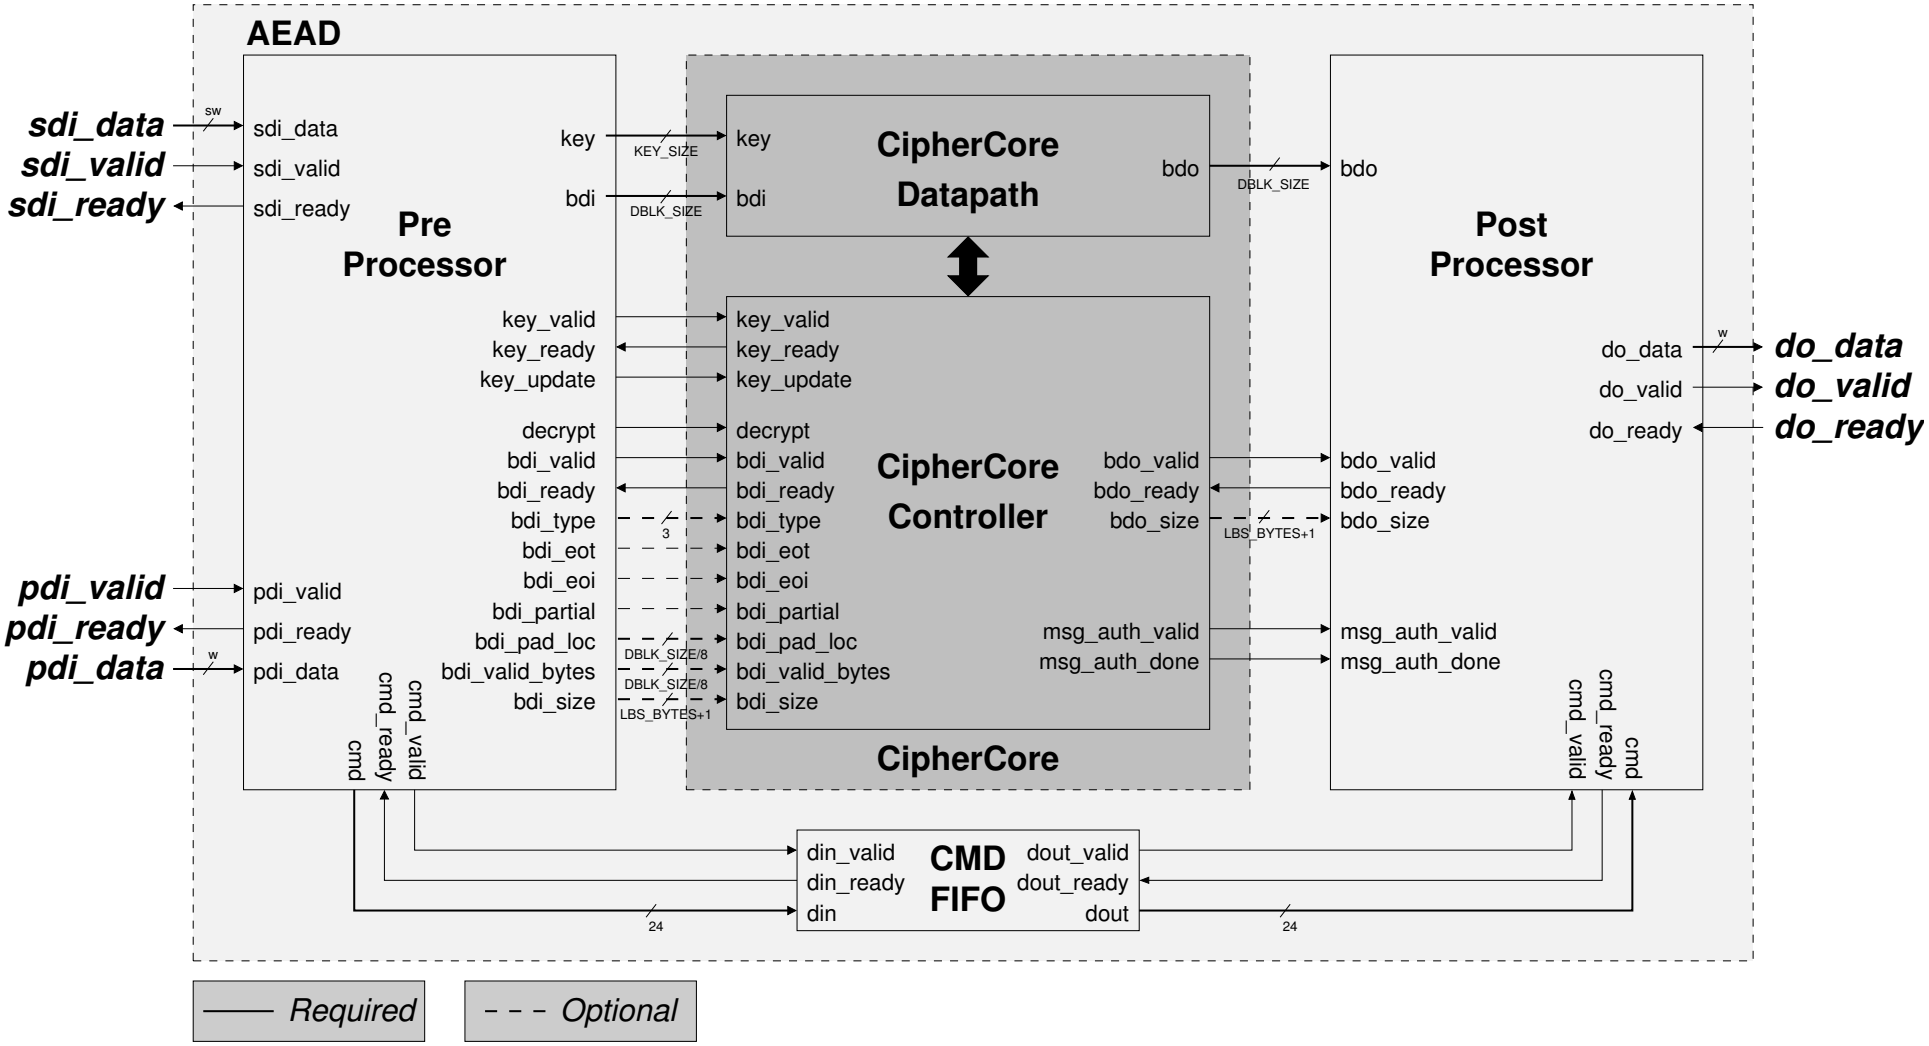

Development Package

May. 12, 2016 - present

a. VHDL code of a generic PreProcessor, PostProcessor, and CMD FIFO, common for all Round 2 Candidates (src_rtl)
b. Universal testbench common for all Round 2 candidates (AEAD_TB)
c. Python app used to automatically generate test vectors (aeadtvgen)
d. Six reference high-speed implementations of Dummy authenticated ciphers (dummyN)

The API Compliant Code Development

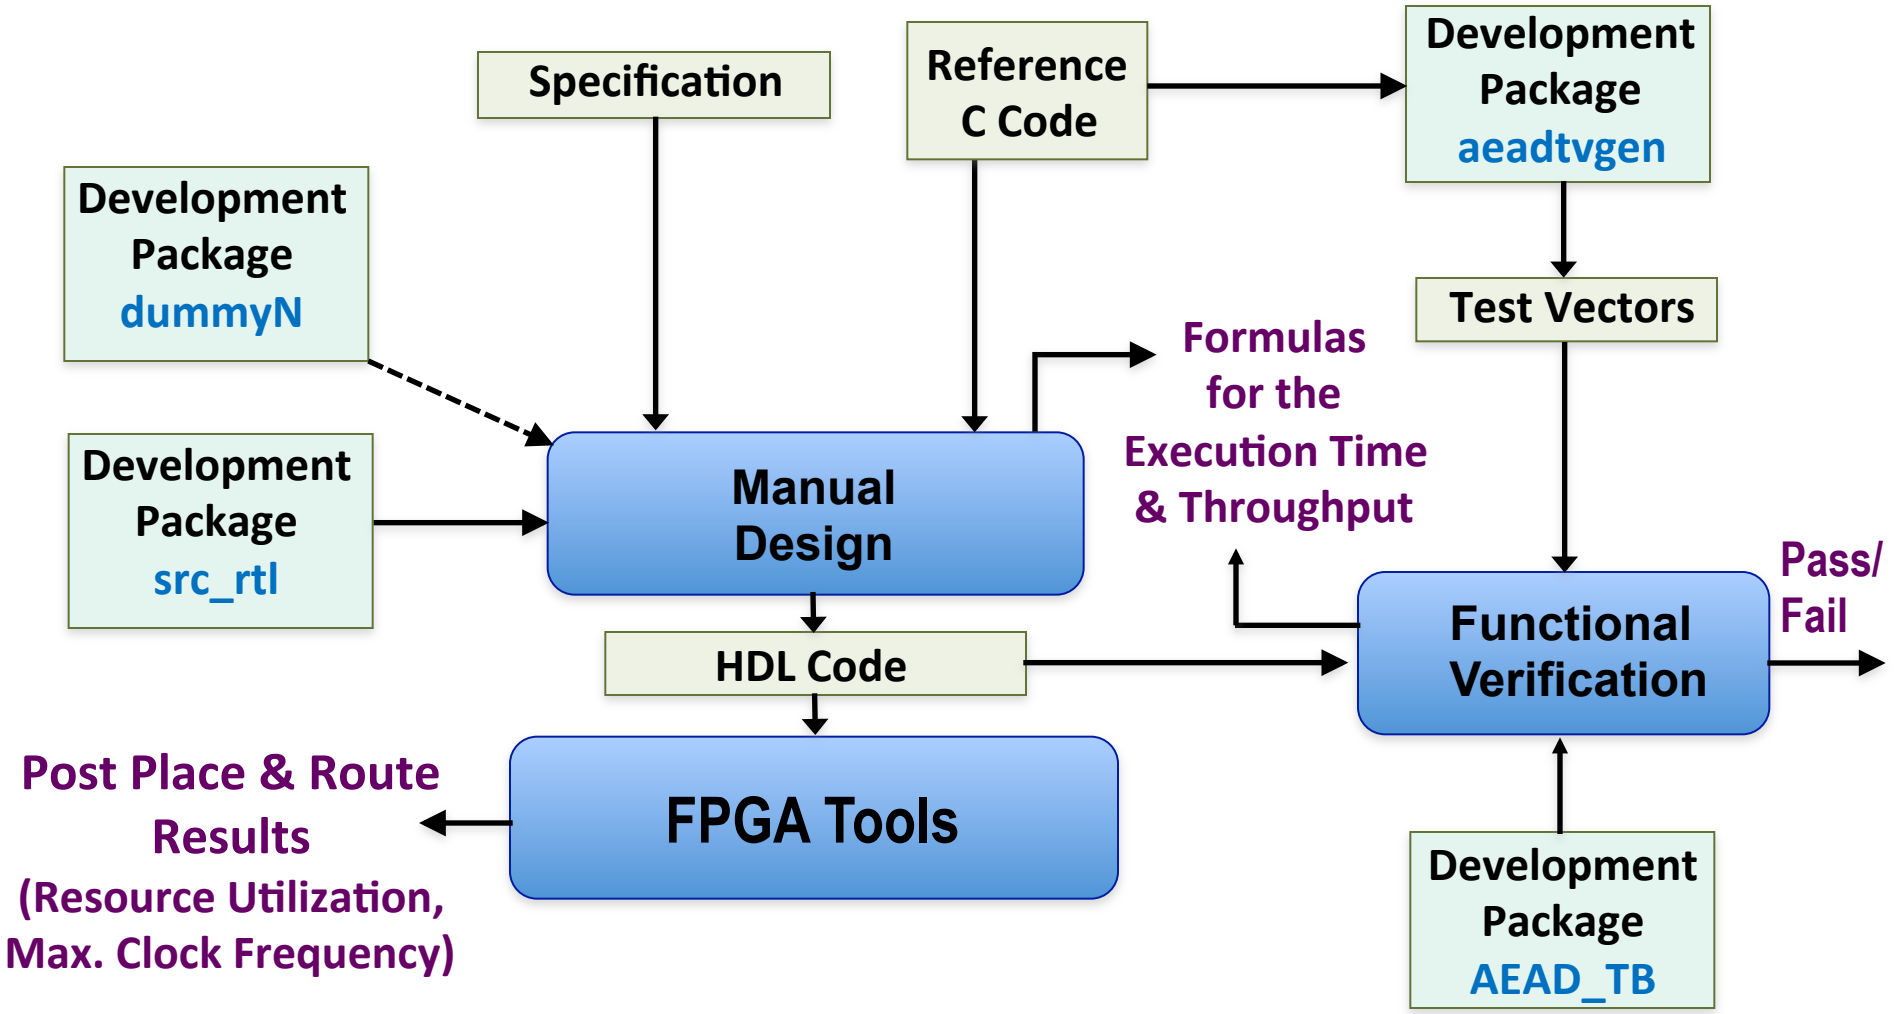

Overview of Submitted Designs

Submitters

1. CCRG NTU (Nanyang Technological University) Singapore – ACORN, AEGIS, JAMBU, & MORUS
2. CLOC-SILC Team, Japan – CLOC & SILC
3. Ketje-Keyak Team – Ketje & Keyak
4. Lab Hubert Curien, St. Etienne, France – ELmD & TriviA-ck
5. Axel Y. Poschmann and Marc Stöttinger – Deoxys & Joltik
6. NEC Japan – AES-OTR
7. IAIK TU Graz, Austria – Ascon
8. DS Radboud University Nijmegen, Netherlands – HS1-SIV
9. IIS ETH Zurich, Switzerland – NORX
10. Pi-Cipher Team – Pi-Cipher
11. EmSec RUB, Germany – POET
12. CG UCL, INRIA – SCREAM
13. Shanghai Jiao Tong University, China – SHELL

Total: 19 Candidate Families

Submitters - GMU Benchmarking Team

"Ice" Homsirikamol
AES-GCM, AEZ,
Ascon, Deoxys,
HS1-SIV, ICEPOLE,
Joltik, NORX, OCB,
PAEQ, Pi-Cipher, STRIBOB

Will Diehl
Minalpher
OMD
POET
SCREAM

Ahmed
Ferozpuri
PRIMATEs-
GIBBON &
HANUMAN,
PAEQ

Farnoud
Farahmand
AES-COPA
CLOC

Mike X.
Lyons
TriviA-ck

Total: 19 Candidate Families + AES-GCM

Variant vs. Architecture

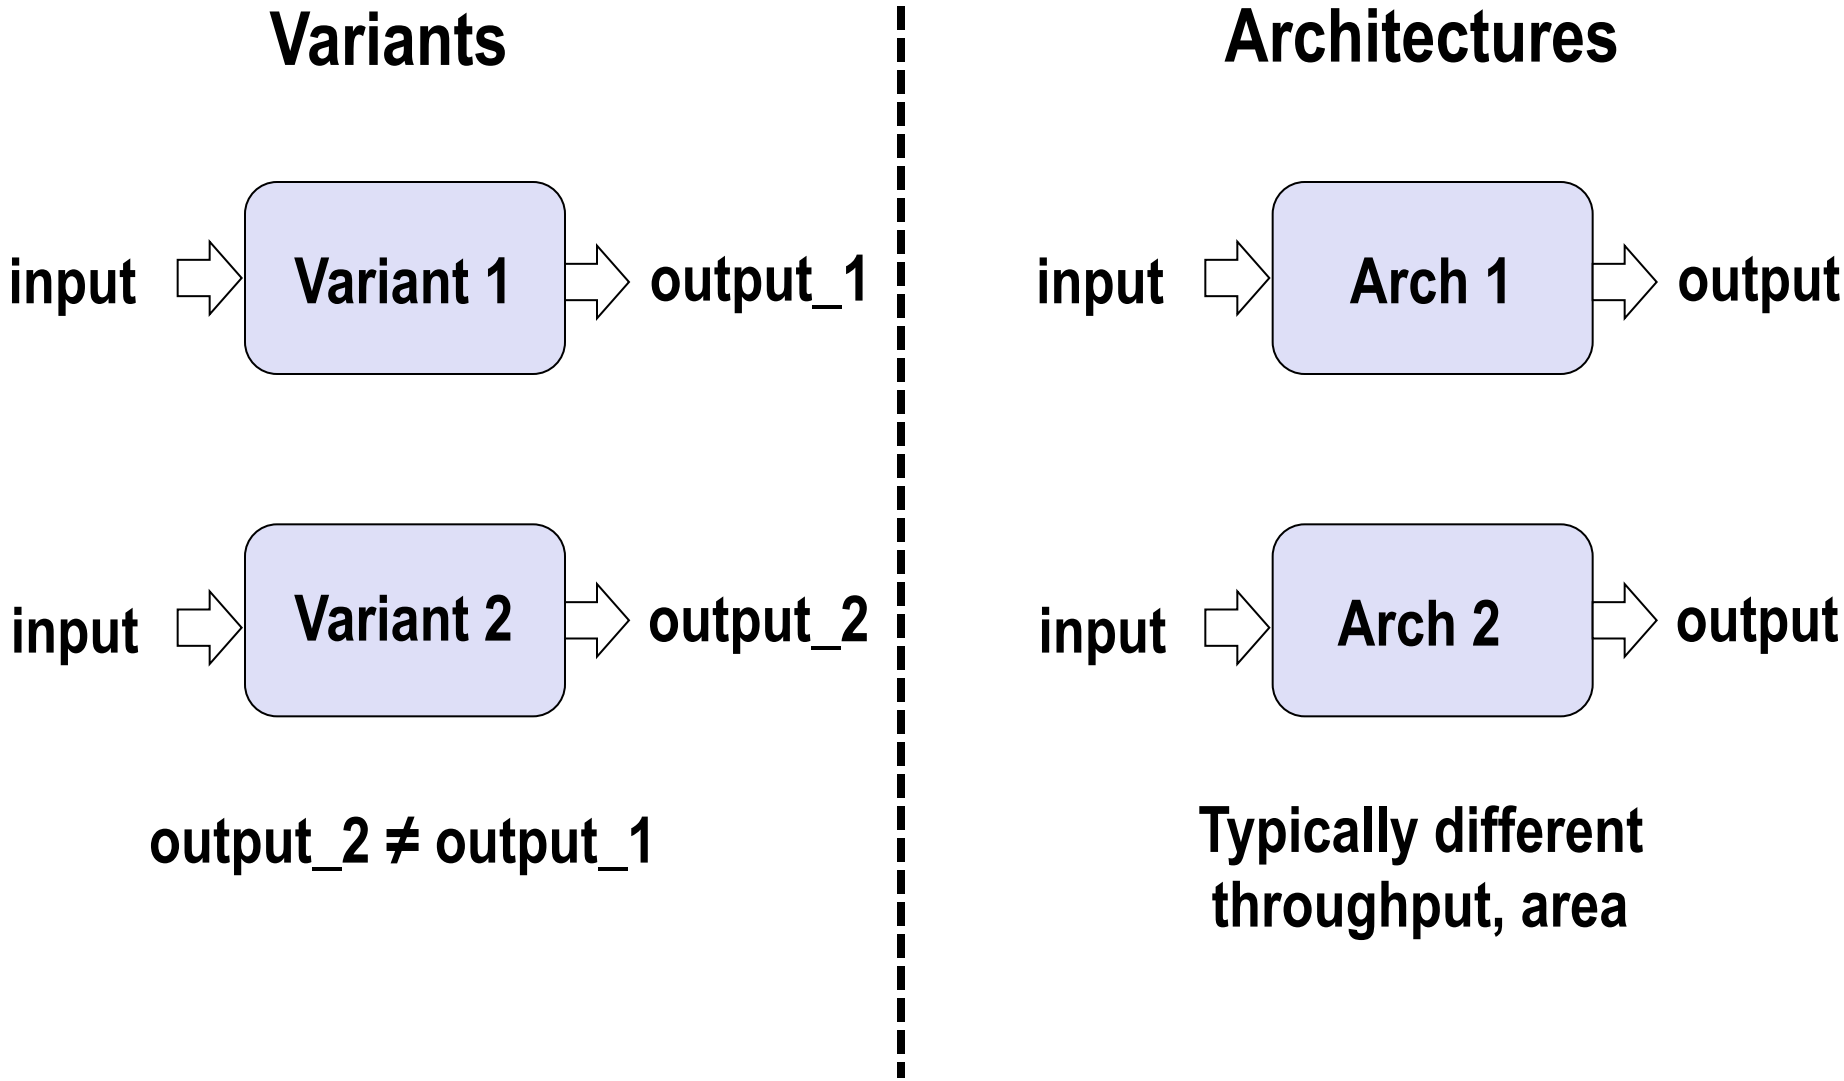

Round 2 Statistics

- **43 hardware design packages**
- **75 variant-architecture pairs**
- Covering the majority of primary variants of

  28 out of 29 Round 2 Candidate Families (all except Tiaoxin)
- High-speed implementation of **AES-GCM** (baseline)

The biggest and the earliest hardware benchmarking effort in the history of cryptographic competitions

Summary of Submitted Designs

- 2 Compliant designs + 1 Non-Compliant Design
  - 1: TriviA-ck
- 2 Compliant designs
  - 3: Ascon, CLOC, Minalpher
- 1 Compliant Design + 1 Non-Compliant Design
  - 8: Deoxys, ELmD, HS1-SIV, Joltik, NORX, Pi-Cipher, POET, SCREAM
- 1 Compliant Design
  - 17: ACORN, AEGIS, AES-COPA, AES-JAMBU, AES-OTR, AEZ, ICEPOLE, Ketje, Keyak, MORUS, OCB, OMD, PAEQ, PRIMATEs-GIBBON, HANUMAN, SHELL, SILC, STRIBOB
- No Designs
  - 1: Tiaoxin

Non Compliant Designs

Algorithm (Target)	Hardware designers	No decryption	Full-block width interface	No support for CAESAR API Protocol	Wrapper required
Deoxys & Joltik (ASIC)	Axel Y. Poschmann & Marc Stöttinger	X	X	X	
POET (ASIC, FPGA)	Amir Moradi		X	X	
SCREAM (ASIC, FPGA)	Lubos Gaspar & Stephanie Kerckhof		X	X	
NORX (ASIC)	Michael Muehl-berghuber		X	X	X

Partial Compliance

Keyak (by the Ketje-Keyak Team)

- Compliance criteria:
  - supported maximum size for AD should be 2^{32}-1 bytes
- Implementation:
  - supported maximum size for AD is **24 bytes**

In the Motorist mode:
metadata (AD) is input together with the plaintext and possibly in input blocks after it

- **Feature unique for Keyak**
- **No plug-in replacement for AES-GCM**

Architectures

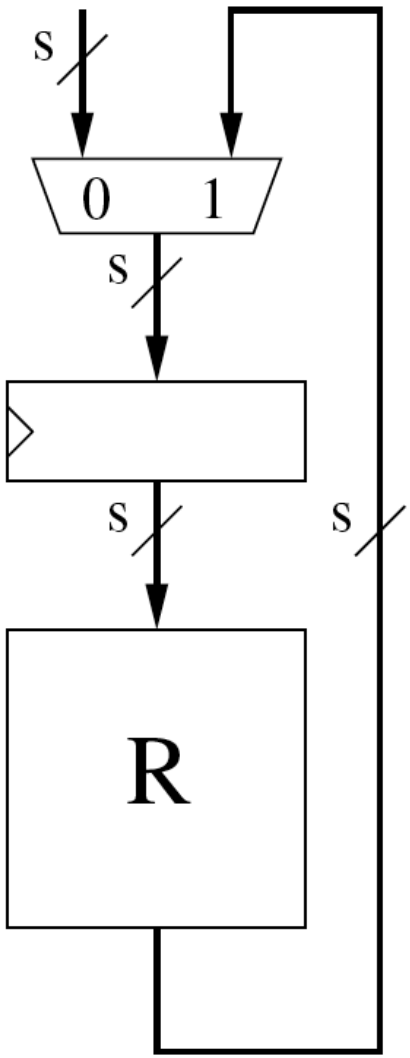

- Majority of algorithms have designs based on

  Basic Iterative Architecture

  - **One round per clock cycle**
  - **Straightforward**
  - **Easy to describe in VHDL/Verilog**
  - **Best or close to best throughput/area**
  - **Hard to optimize**

Other Architectures:

- Lightweight: ACORN
- Folded: HS1-SIV, Pi-Cipher
- Unrolled (extra): Ascon, SCREAM
- With Speculative Precomputation: Deoxys

Key sizes

- Majority of implemented ciphers support **128-bit keys only**

  Exceptions:
  - AES-JAMBU, Ketje: 96
  - AEZ: 384
  - PRIMATEs: 80 & 120
  - STRIBOB: 192
  - Joltik: 64 & 128
  - Pi-Cipher: 96, 128, 256
  - Deoxys, NORX: 128 & 256

Possible allowed key ranges:

$|K| \geq 96$

- **covers all families**
- **excludes variants with 64 and 80-bit keys**

$|K| \geq 120$

- covers all families except AES-JAMBU and Ketje
- covers stronger variants of PRIMATEs
- excludes lightweight variants

PDI & DO Ports Width, w

- The CAESAR API Minimum Compliance Criteria allow
  - **High-speed: $32 \leq w \leq 256$**
  - **Lightweight: $w = 8, 16, 32$**

- Majority of the API compliant implementations support **w=32 or 64 only**

  Exceptions:
  - ACORN: 8 & 32
  - PRIMATEs: 40
  - HS1-SIV: 128
  - NORX, Pi-Cipher: 128 & 256
  - AEGIS, ICEPOLE, MORUS: 256

Benchmarking Methodology

FPGA Families & Devices Used for Benchmarking

High-Performance FPGA Families used for benchmarking of
All Round 2 Candidates & AES-GCM

- Xilinx Virtex-6: xc6vlx240tff1156-3
- Xilinx Virtex-7: xc7vx485tffg1761-3
- Altera Stratix IV: ep4se530h35c2
- Altera Stratix V: 5sgxea7k2f40c1

Low-Cost FPGA Families used for benchmarking of
10 Candidates with the Smallest Area in High-Performance Benchmarking:

- Xilinx Spartan-6: xc6slx16csg324-3
- Xilinx Artix-7: xc7a100tcsg324-3
- Altera Cyclone IV: EP4CE22F17C6
- Altera Cyclone V: 5CEBA4F23C7

RTL Benchmarking

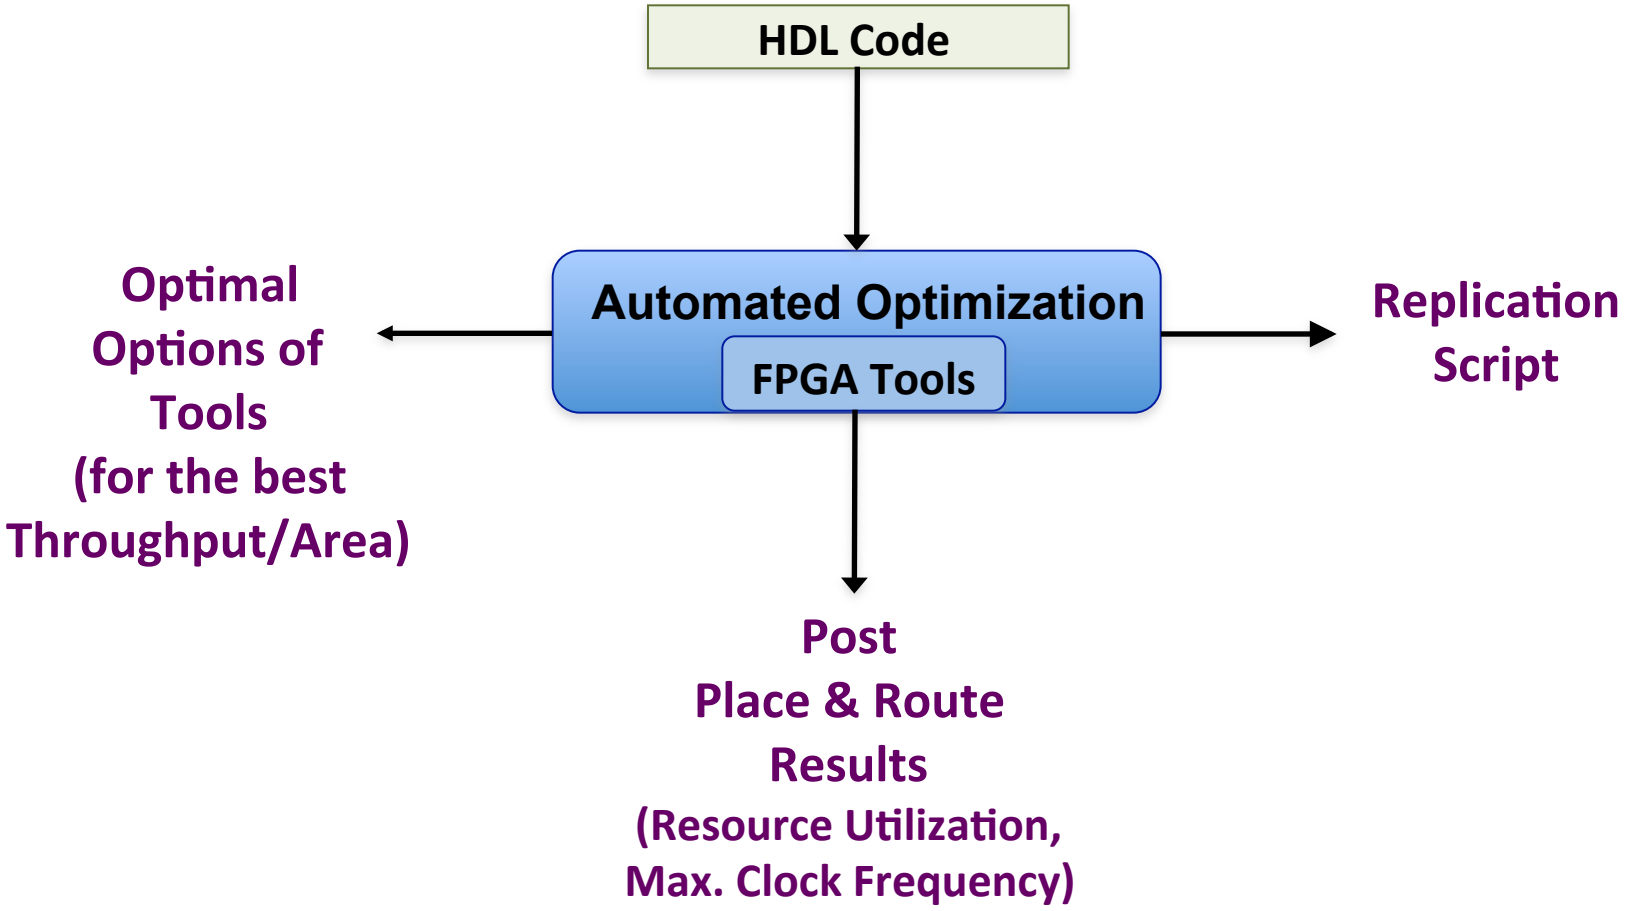

FPGA Tools (1)

For Benchmarking Targeting Xilinx FPGAs (other than Virtex 7):

Target FPGAs:	Virtex-6, Spartan 6, Artix 7
Synthesis Tool:	Xilinx XST 14.7
Implementation Tool:	Xilinx ISE 14.7
Automated Optimization:	ATHENa

For Benchmarking Targeting Altera FPGAs:

Target FPGAs:	Stratix IV, Stratix V, Cyclone IV, Cyclone V
Synthesis Tool:	Quartus Prime 16.0.0
Implementation Tool:	Quartus Prime 16.0.0
Automated Optimization:	ATHENa

FPGA Tools (2)

For Benchmarking Targeting Xilinx Virtex 7 FPGAs:

Target FPGAs:	**Virtex-7**
Synthesis Tool:	**Xilinx Vivado 2015.1**
Implementation Tool:	**Xilinx Vivado 2015.1**
Automated Optimization:	**25 Default Strategies of Vivado**

Results

Virtex-6

Results for Virtex 6 – Throughput vs. Area
Linear Scale

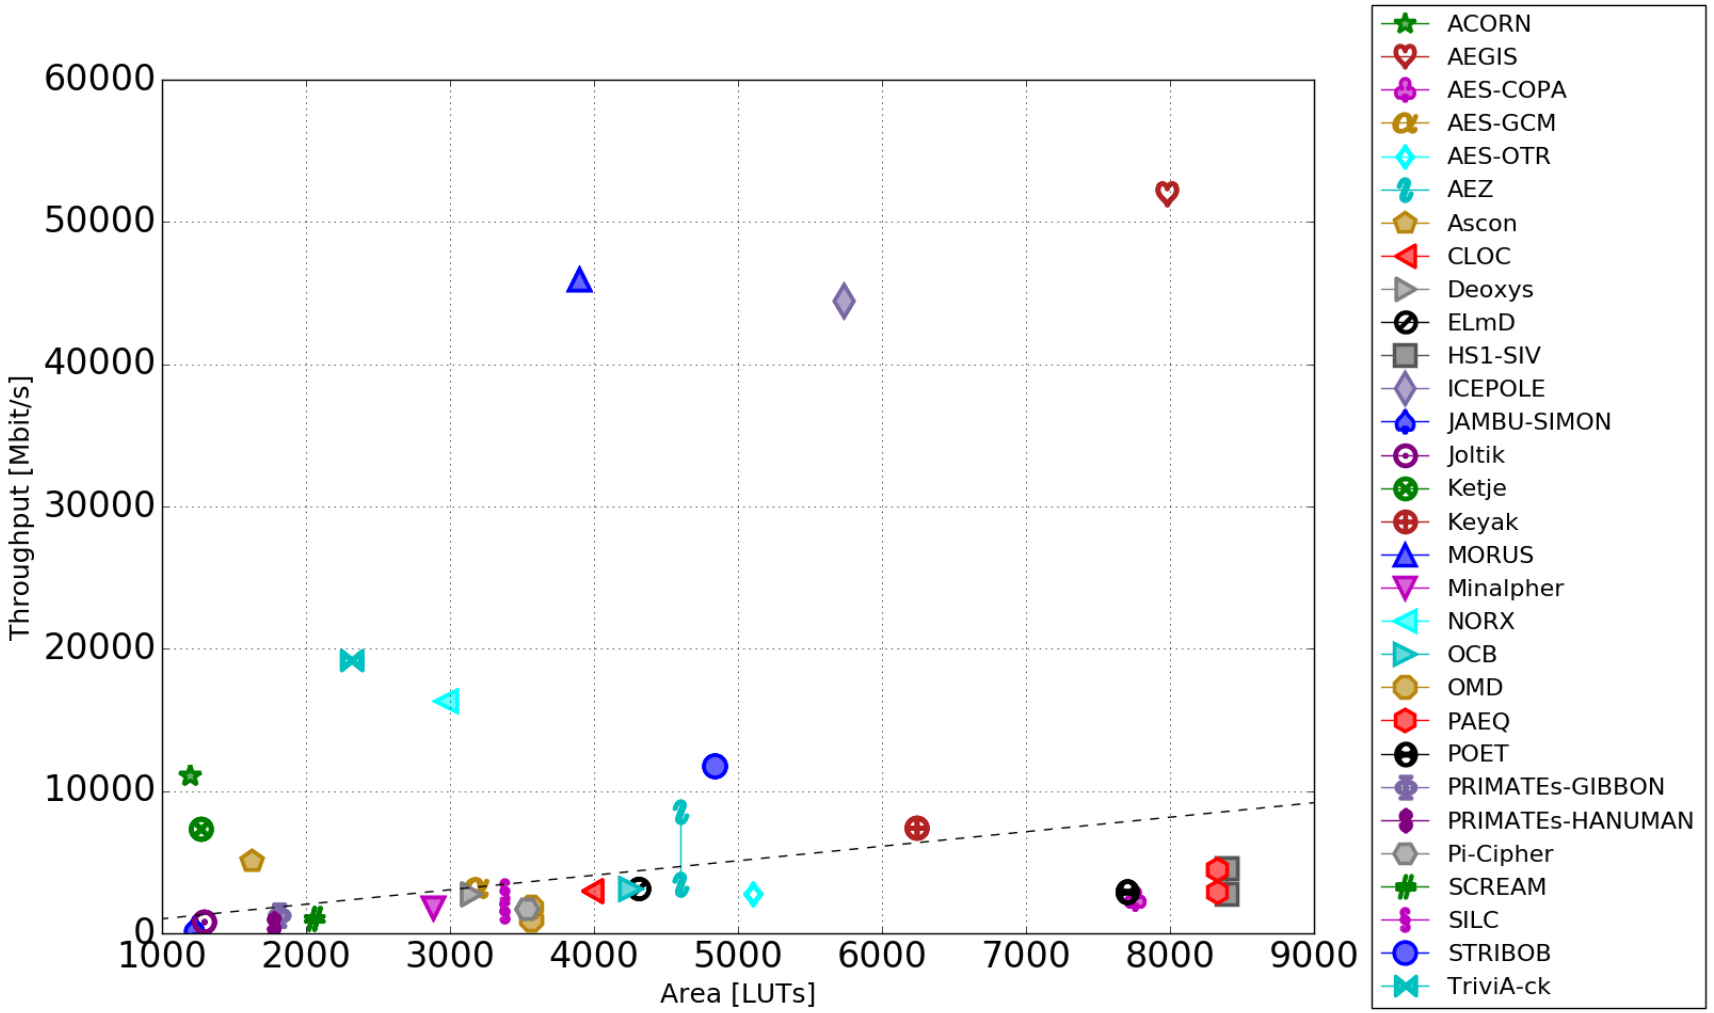

Results for Virtex 6 – Throughput vs. Area Logarithmic Scale

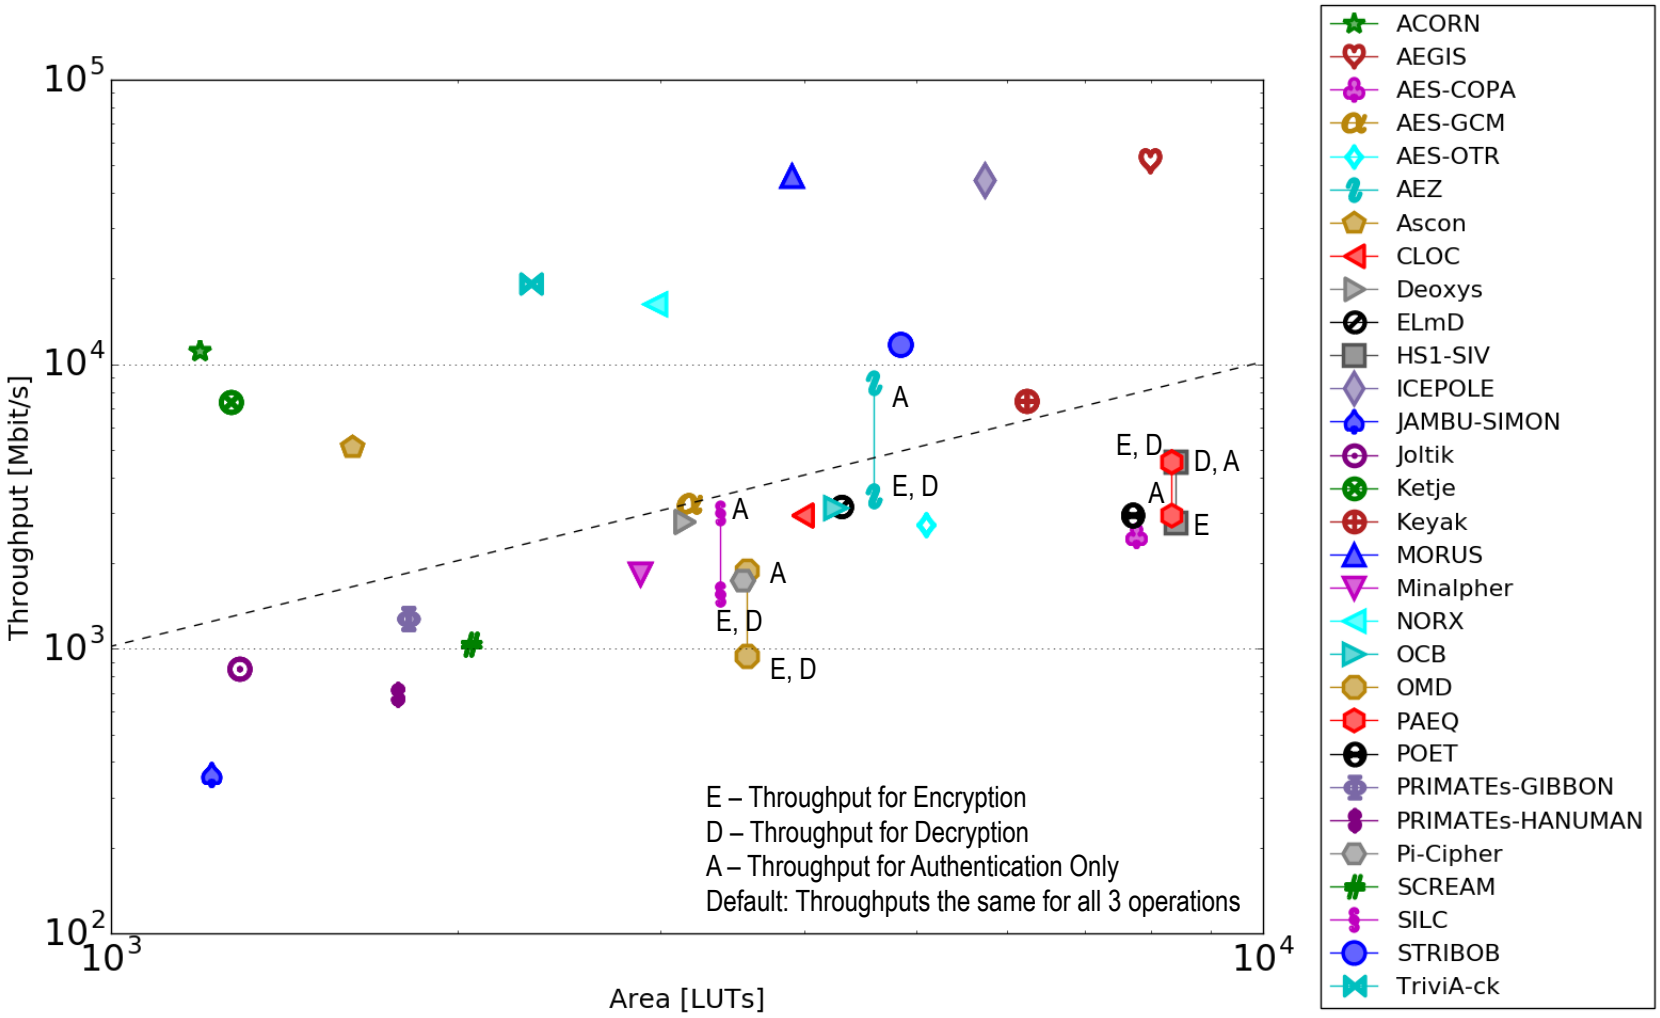

Relative Throughput/Area in Virtex 6 vs. AES-GCM

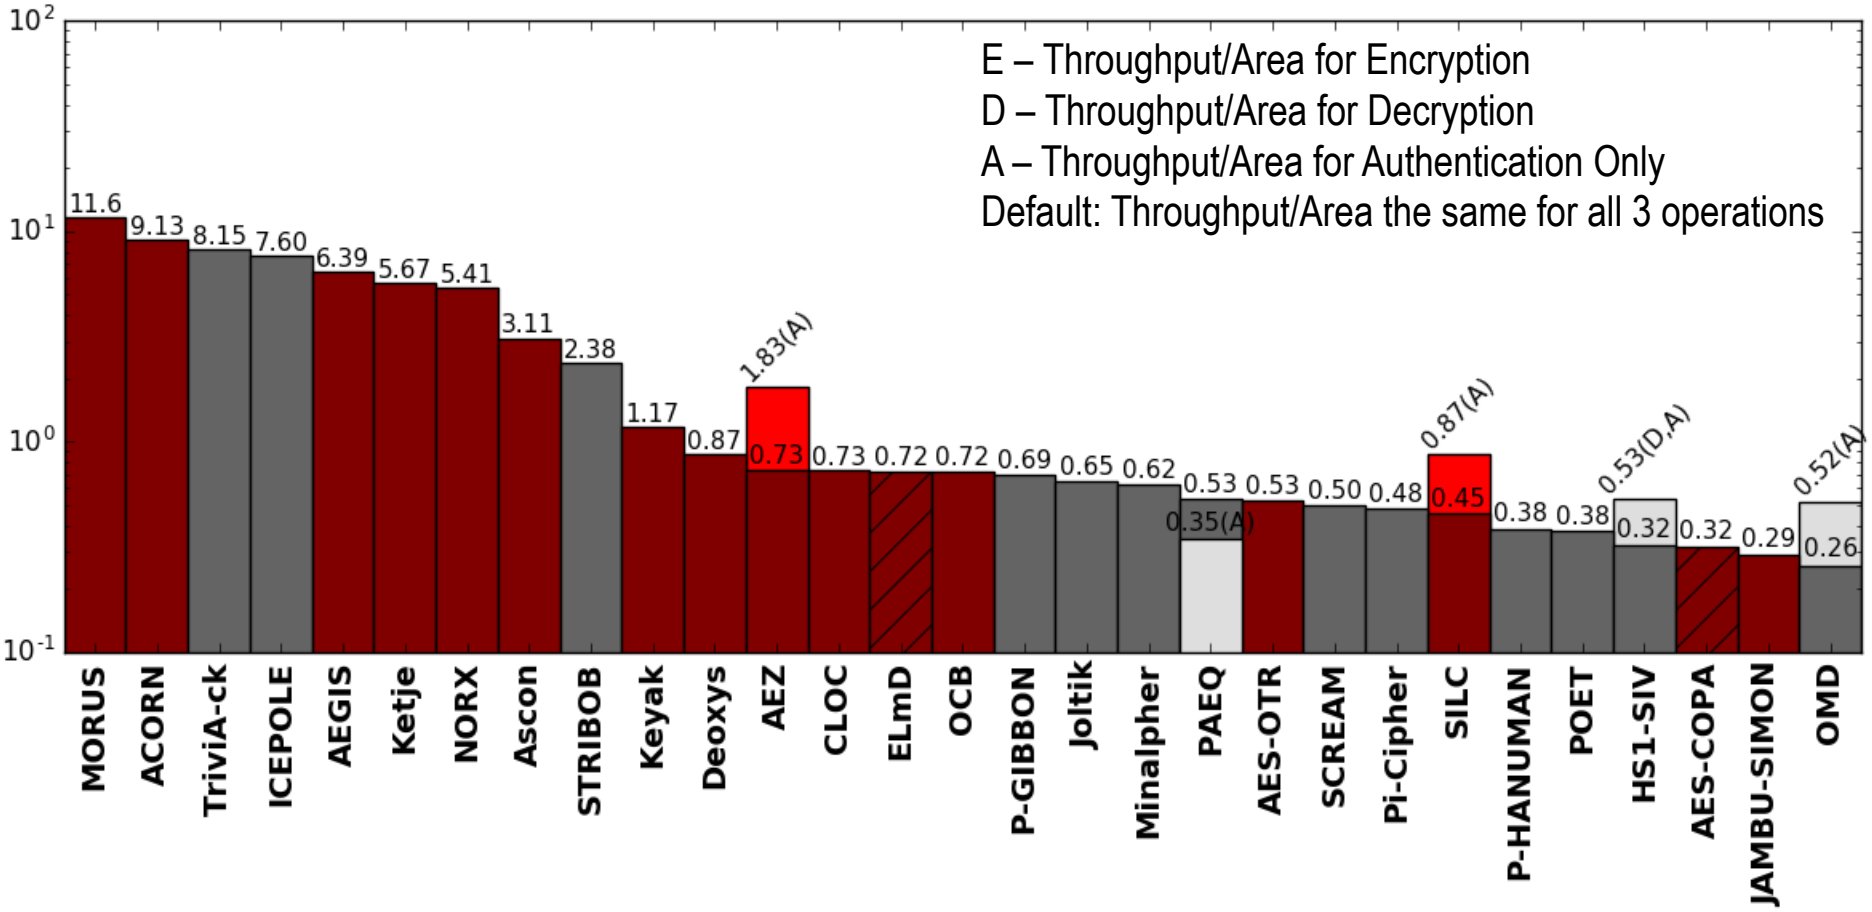

Throughput/Area of AES-GCM = 1.020 (Mbit/s)/LUTs

Relative Throughput in Virtex 6
Ratio of a given Cipher Throughput/Throughput of AES-GCM

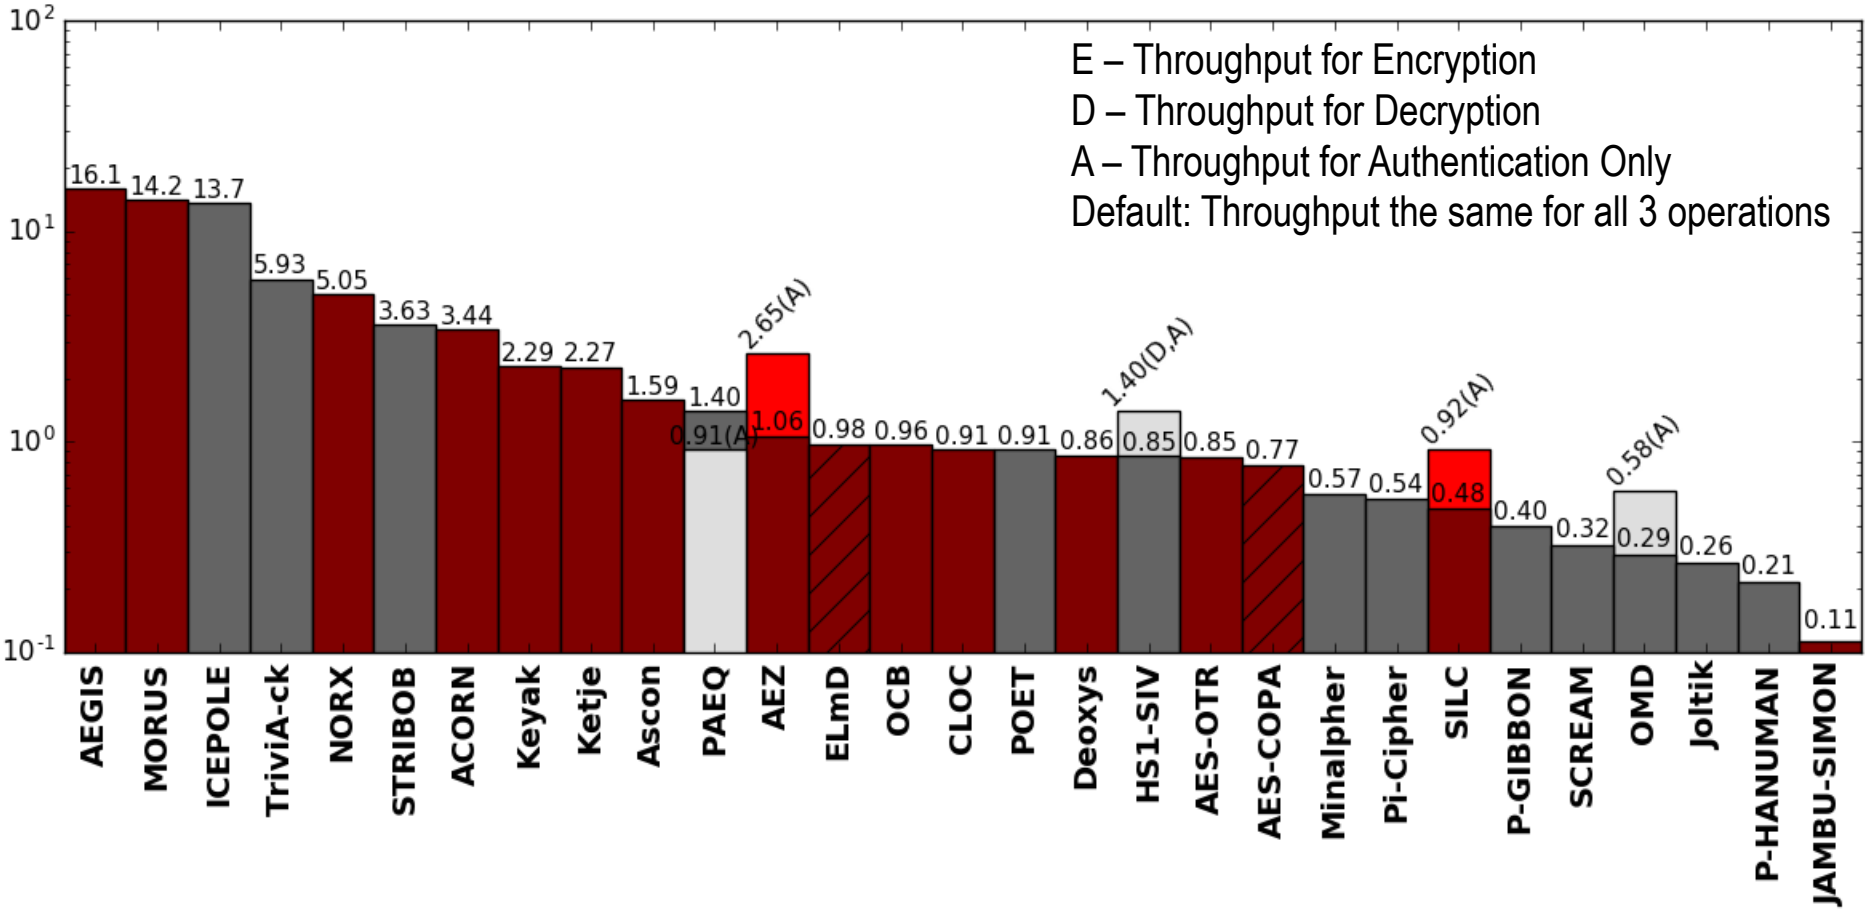

Throughput of AES-GCM = 3239 Mbit/s

Relative Area (#LUTs) in Virtex 6
Ratio of a given Cipher Area/Area of AES-GCM

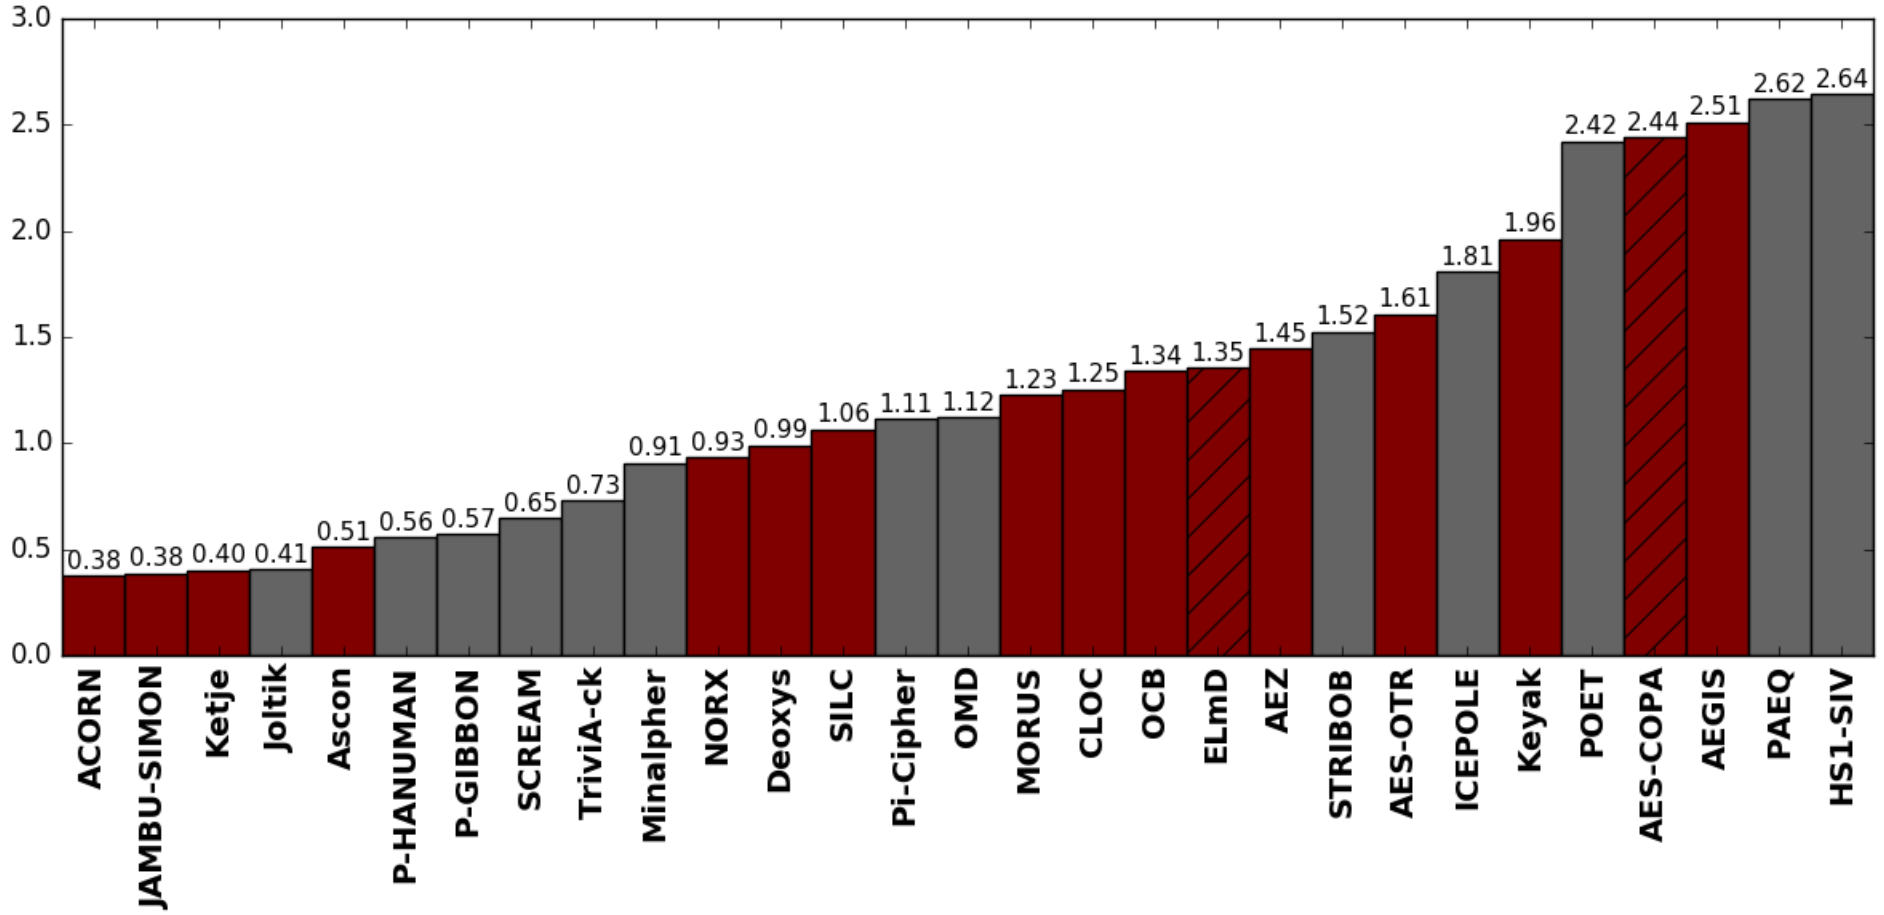

Area of AES-GCM = 3175 LUTs

ATHENa Database of Results

ATHENa Database of Results

- Available at
  http://cryptography.gmu.edu/athena

- Developed by John Pham, a Master's-level student of Jens-Peter Kaps as a part of the SHA-3 Hardware Benchmarking project, 2010-2012, (sponsored by NIST)

- In June 2015 extended to support Authenticated Ciphers

One Stop Website

https://cryptography.gmu.edu/athena/index.php?id=CAESAR
OR
https://cryptography.gmu.edu/athena
and click on Download

- VHDL/Verilog Code of CAESAR Candidates: Summary I
- VHDL/Verilog Code of CAESAR Candidates: Summary II
- ATHENa Database of Results: Rankings View
- ATHENa Database of Results: Table View
- Benchmarking of Round 2 CAESAR Candidates in Hardware: Methodology, Designs & Results
- GMU Implementations of Authenticated Ciphers and Their Building Blocks
- CAESAR Hardware API v1.0

Round 3
Benchmarking
Goals & Timeline

Round 3 Candidates Outperforming AES-GCM

High-Speed Implementations (4 FPGA families)

Throughput/Area:

1. ACORN
2. AEGIS
3. Ascon
4. Ketje
5. Keyak
6. MORUS
7. NORX

Throughput:

1. ACORN
2. AEGIS
3. Ascon
4. Ketje
5. Keyak
6. MORUS
7. NORX

Alphabetical Order

R3 Candidates – Relative Throughput/Area - Virtex 6

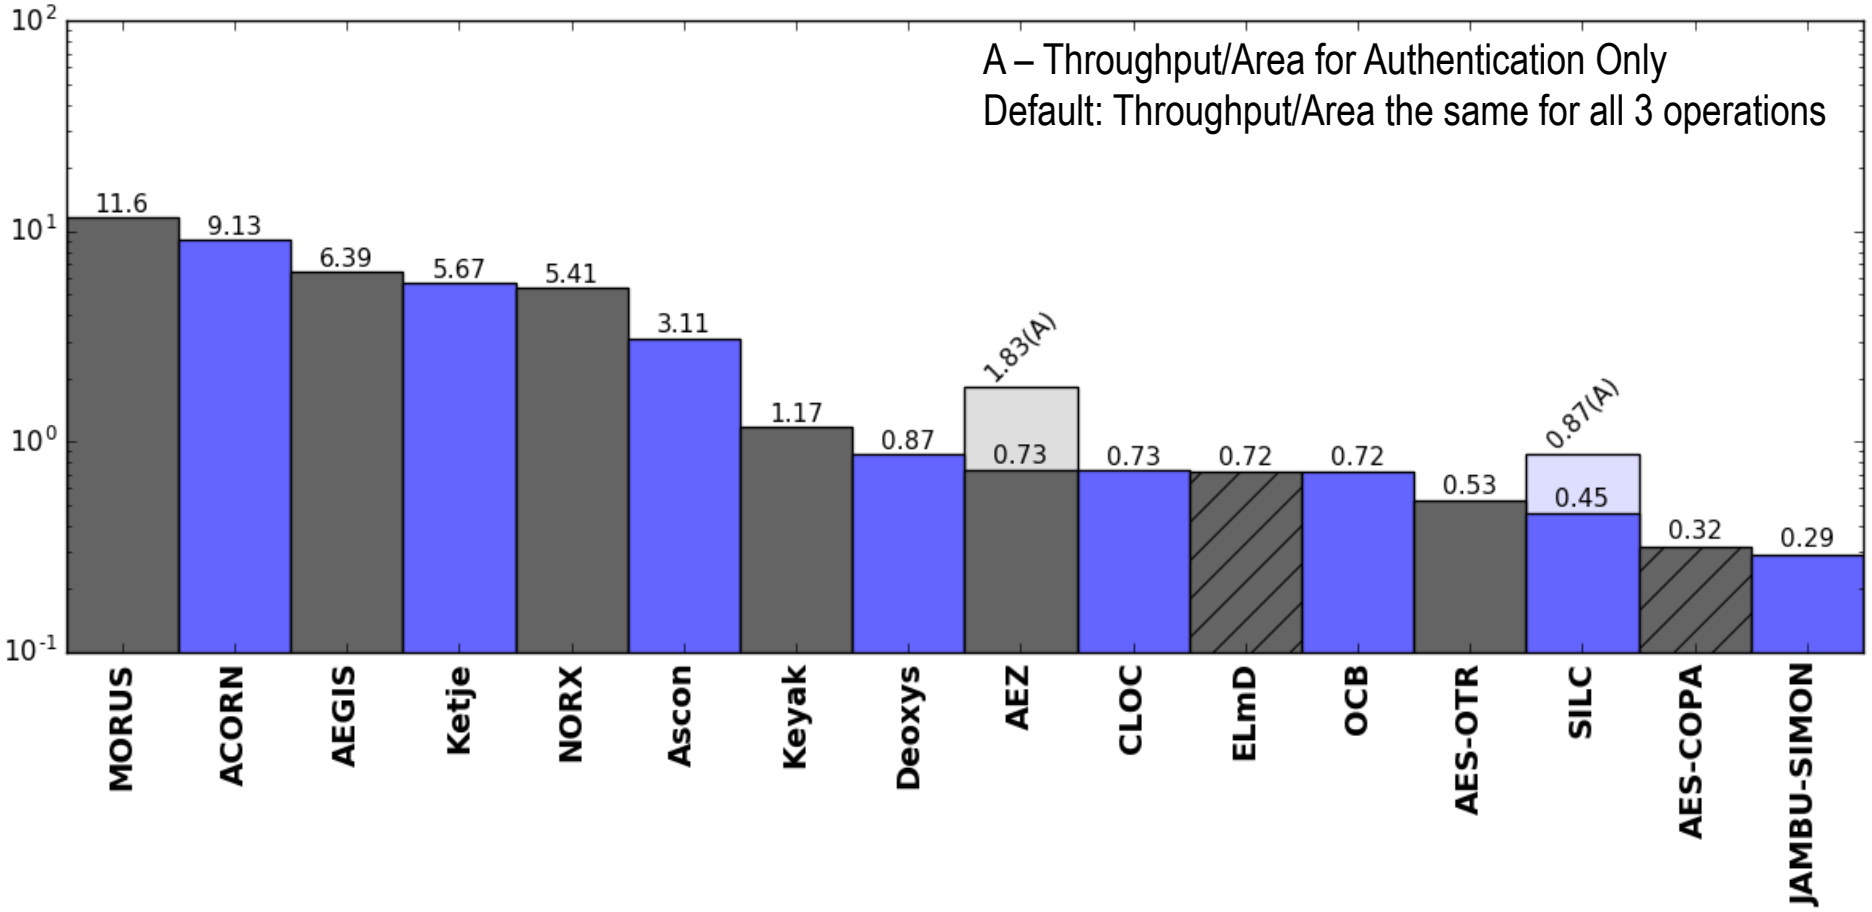

Throughput/Area of AES-GCM = 1.020 (Mbit/s)/LUTs

R3 Candidates – Relative Throughput - Virtex 6

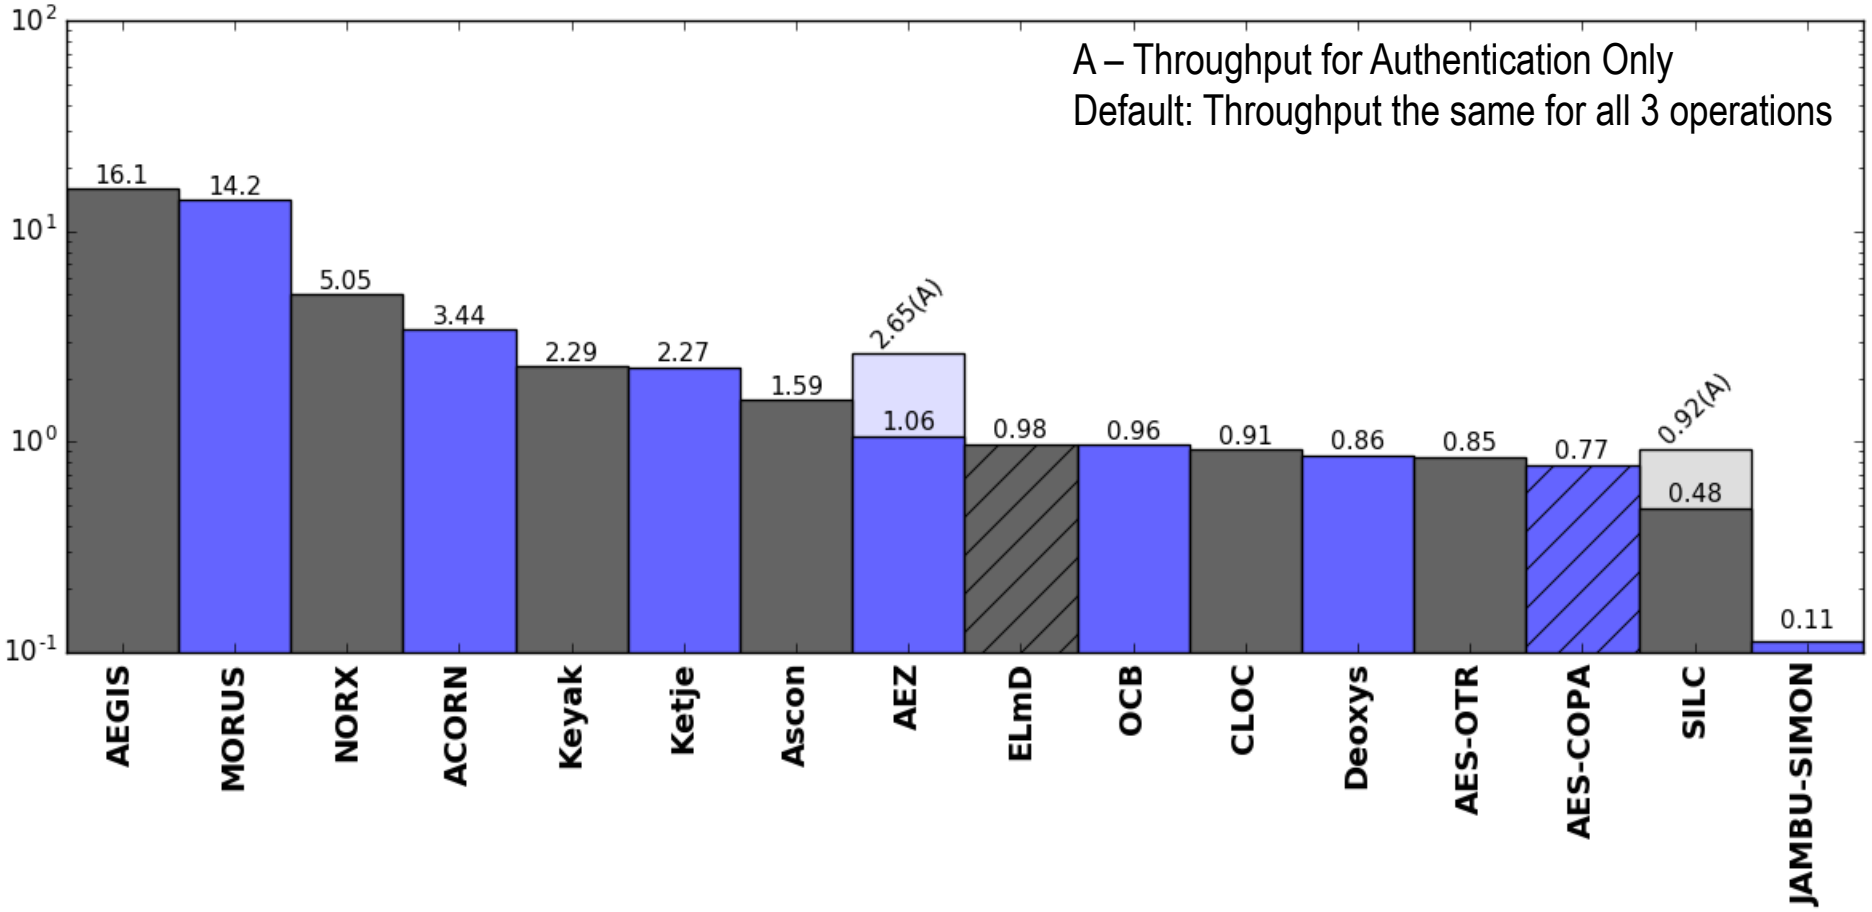

Throughput of AES-GCM = 3239 Mbit/s

Round 3 Benchmarking Goals

I. Lightweight Implementations, benchmarked for area, throughput/area, power, energy/bit

1. ACORN
2. Ascon
3. CLOC (TWINE-80, AES-128
4. JAMBU (SIMON, AES)
5. Ketje
6. SILC (PRESENT-80, LED-80, AES-128)
7. Others (AES-OTR, COLM, Deoxys, Keyak, MORUS)?

II. Natural resistance to side-channel attacks & the cost of countermeasures

Possibly a subject of the next DPA Contest ?

Round 3 Benchmarking Goals

III. ASIC Benchmarking

- High-speed implementations
- Lightweight implementations
- Implementations of two-pass algorithms (effect of external memory)
- Side-channel resistance

IV. High-speed architectures supporting multiple messages processed in parallel

- Multi-message pipelining
- Extensions to API required

Round 3 Benchmarking Goals

V. Investigating Throughputs vs. Area Trade-offs (flexibility, wide range of applications)

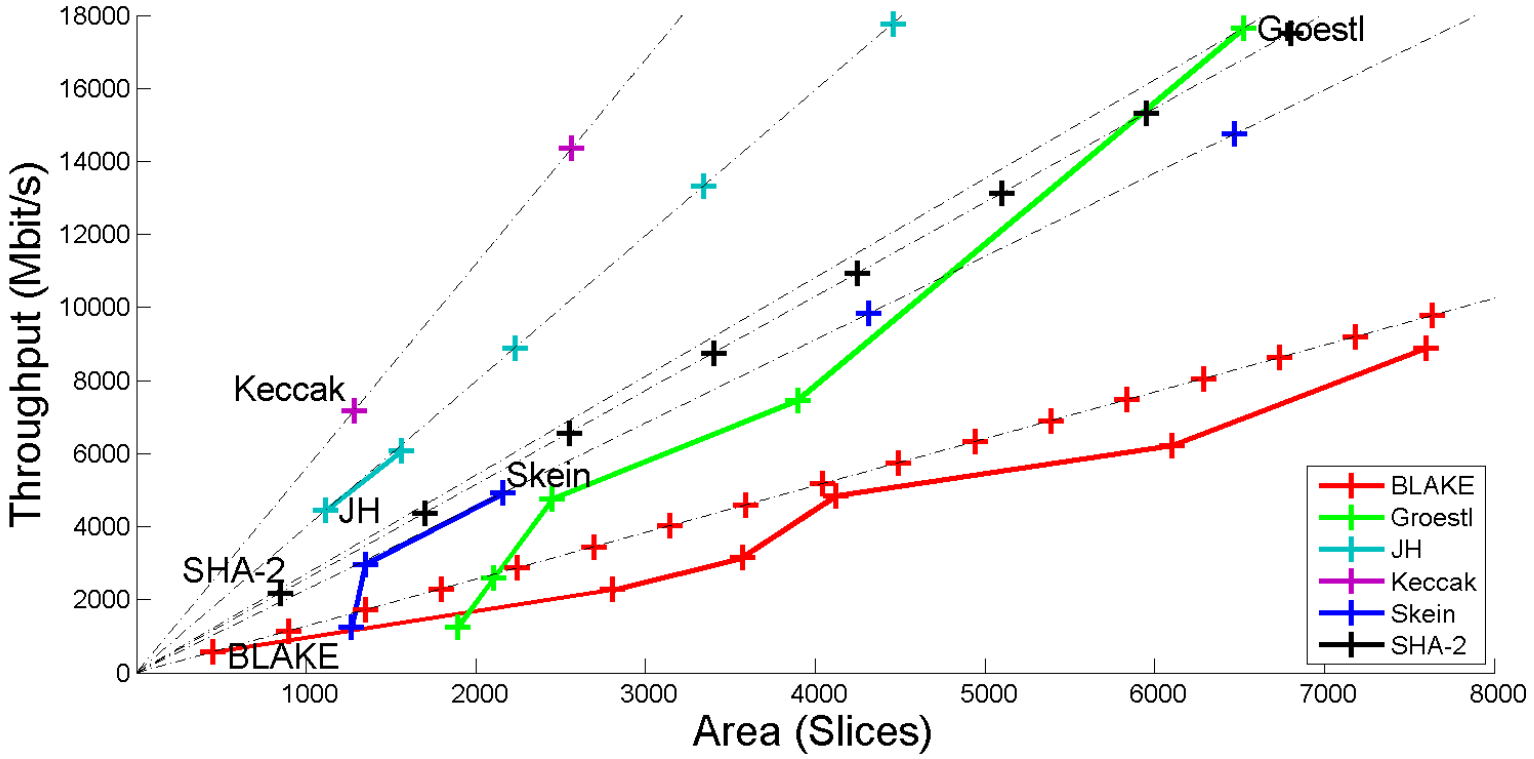

Possible Architectures: folded, unrolled, with inner-round pipelining, etc.

Round 3 Benchmarking Goals

VI. Extensions Common for all Authenticated Ciphers

- buffering of decrypted data before authentication
- merging Npub, AD, Ciphertext, and Tag after decryption
- word width conversion (for communication between implementations with different PDI/SDI/DO widths)

VI. Experimental Setups

- power/energy measurements
- communication & control overhead of a hardware accelerator
- operating system overhead
- CAESAR API validation taking into account the most popular Bus Interfaces, such as AXI4 and PCIe

Round 3 Benchmarking Timeline

Requests for changes in the CAESAR API:

October 31, 2016

Round 3 VHDL/Verilog:

At least **two months before** the
announcement of finalists

Independent Benchmarking Efforts (ASIC, Side-channel, etc.):

Early declarations and guidelines
for designers strongly encouraged

Conclusions

- The biggest and the earliest hardware benchmarking effort in the history of cryptographic competitions
  - 14 hardware designer groups
  - 28 candidate families
  - 75 variant-architecture pairs
- Key new features:
  - Standard API
  - Implementer's Guide and Development Package
  - Algorithm designers requested to submit HDL code (possibly designed by other teams)
- Modest but noticeable influence on the Round 3 selection

Possible Improvements

- Faster adoption of the submitted proposals (e.g., API) by the CAESAR Committee
- More realistic and relaxed deadlines
- Clear indication of the influence of hardware benchmarking on the final decision
  - Avoiding mixed signals:
    - ➢ "reference" hardware implementation
    - ➢ advancing candidates without VHDL/Verilog code
- Early collaborations
- More groups involved in various benchmarking efforts (lightweight, ASIC, side-channel)
- Incentives: publication venues, grants, PhD/MS theses

Thank you!

Questions?

Comments?

Suggestions?

ATHENa: http://cryptography.gmu.edu/athena
CERG: http://cryptography.gmu.edu