CERG

Cryptographic Engineering Research Group

cryptography.gmu.edu

FPGA Benchmarking of Round 2 Candidates in the NIST Lightweight Cryptography Standardization Process: Methodology, Metrics, Tools, and Results

Kamyar Mohajerani, Richard Haeussler, Rishub Nagpal,
Farnoud Farahmand, Abubakr Abdulgadir,
Jens-Peter Kaps and Kris Gaj

George Mason University
USA

GMU CERG LWC Benchmarking Team

Kamyar
Mohajerani

Richard
Haeussler

Rishub
Nagpal

Farnoud
Farahmand

Bakry
Abdulgadir

FPGA Benchmarking Goals

- Stimulate the development of hardware implementations that can be fairly compared with each other
- Perform design space exploration of at least selected candidates
- Evaluate and rank candidates from the point of view of their performance in FPGAs
- Identify the best and worst performers in terms of major benchmarking metrics
- Develop optimized code of unprotected implementations to be used as a basis for the development and analysis of protected implementations in Round 3

Previous Similar Benchmarking Efforts

CAESAR – Competition for Authenticated Encryption: Security, Applicability, & Robustness

- HDL code requirement established by the CAESAR Committee in the middle of Round 2 in May 2016

- CAESAR Round 2 : 2015-2016
  - 14 Hardware Design Teams
  - 28 out of 29 candidates implemented

- CAESAR Round 3 : 2016-2017
  - 10 Hardware Design Teams
  - 15 out of 15 candidates implemented

LWC Hardware API proposed by GMU, TUM, & VT

1. Minimum Compliance Criteria

- Supported operations
- Permitted input sizes
- Decrypted plaintext release
- Permitted data port widths

etc.

2. Interface

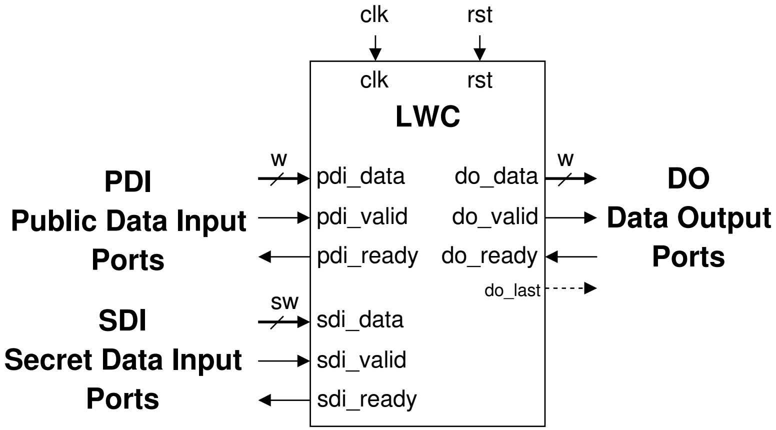

3. Communication Protocol

instruction = ACTKEY
instruction = ENC
seg_0_header
seg_0 = Npub
seg_1_header
seg_1 = AD
seg_2_header
seg_2 = Plaintext

4. Timing Characteristics

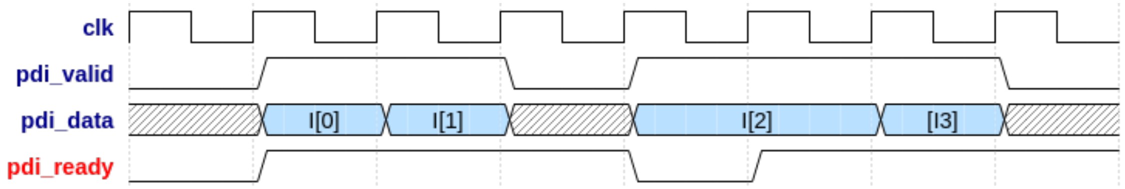

Based on the CAESAR API. Stable since October 2019.

LWC Hardware Development Package by GMU, TUM, & VT

Helpful for all designers

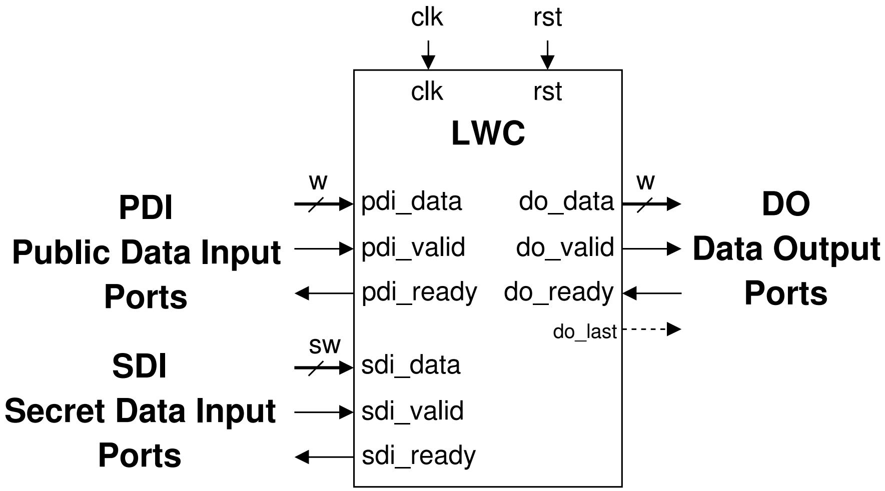

- Universal testbench (LWC_TB)
- Test vector generator (cryptotvgen)

Helpful for designers following the recommended design flow

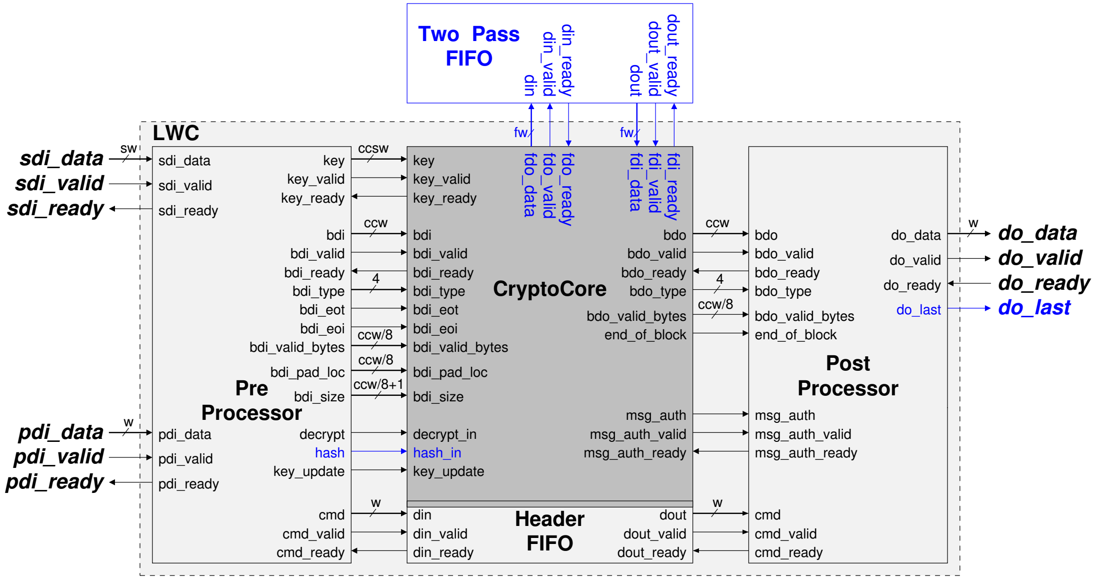

- VHDL code of a generic PreProcessor, PostProcessor, and Header FIFO
- CryptoCore of a dummy authenticated cipher with hash functionality

Choice of FPGA Platforms and Tools

1. Representing widely used low-cost, low-power FPGA families
2. Capable of holding SCA-protected designs

   (possibly using up to 4 times more resources than unprotected designs)
3. Supported by free versions of state-of-the-art industry tools

FPGA Platforms and Tools

Xilinx:

Artix-7 : xc7a12tcsg325-3

8,000 LUTs 16,000 FFs 40 18Kbit BRAMs 40 DSPs 150 I/Os

Intel:

Cyclone 10 LP : 10CL016-YF484C6

15,408 LEs 15,408 FFs 56 M9K blocks 56 MULs 162 I/Os

Lattice Semiconductor:

ECP5 : LFE5U-25F-6BG381C

24,000 LUTs 24,000 FFs 56 18Kbit blocks 28 MULs 197 I/Os.

Xilinx Artix-7 LUTs have 6 inputs; Cyclone 10 LP and ECP5 LUTs have 4 inputs

Optimization Targets

Maximum Throughput with up to

2000 LUTs, 4000 flip-flops of Artix-7 FPGA. No BRAMs & no DSP units.

Alternative targets:

1. Basic-iterative architecture
2. Architectures most natural for a given authenticated cipher
   a. Folding in block-cipher-based submissions
   b. Generating 2^d bits per clock cycle in stream-cipher-based submissions
3. Maximum Throughput for 1000 LUTs, 2000 flip-flops of Artix-7 FPGA.
   No BRAMs & no DSP units.

Benchmarking Metrics (1)

Metrics obtained from tool reports after placing and routing:

1. Resource utilization

   Number of LUTs (LEs for Cyclone 10LP) and flip-flops, assuming no use of embedded memories (such as BRAMs), DSP units, and embedded multipliers

2. Maximum clock frequency in MHz

   (used only for the calculation of maximum throughput)

Benchmarking Metrics (2)

Metrics calculated based on the execution time measurements (in clock cycles) obtained using functional simulation and the maximum clock frequencies (in MHz):

Throughput in Mbits/s

for the following sizes of inputs

a. long [with Throughput = d · Block size/(Time(N+d blocks)-Time(N blocks))]
b. 1536 bytes
c. 64 bytes
d. 16 bytes.

All throughputs calculated separately for

- authenticated encryption: AD, plaintext, AD+plaintext (sender's side)
- authenticated decryption: AD, ciphertext, AD+ciphertext (receiver's side)
- hashing: hash message (both sides)

Timeline of Round 2 FPGA Benchmarking

Phase 1:

Sep. 1, 2020:	1st submission deadline
Sep. 26, 2020:	Publication of the living report

Phase 2:

Oct. 11, 2020:	2nd submission deadline
Oct. 21, 2020:	Phase 2 updates to the report

Phase 3:

Nov. 9, 2020:	3rd submission deadline
Nov. 30, 2020:	Final version of the report

Summary of Submissions
Phases 1 & 2

27 submissions representing 22 out of 32 candidates (69%)

Candidates with two submissions from two different groups:

Ascon, COMET, Gimli, TinyJAMBU, and Xoodyak

Summary of Submissions
Phases 1 & 2

- George Mason University Cryptographic Engineering Research Group (CERG), USA (6):

  Elephant, PHOTON-Beetle, Pyjamask, Saturnin, TinyJAMBU, and Xoodyak
- Virginia Tech Signatures Analysis Lab, USA (5):

  Ascon, COMET, GIFT-COFB, SCHWAEMM & ESCH, and Spoc
- CINVESTAV-IPN, Mexico (4)

  COMET, ESTATE, LOCUS-AEAD/LOTUS-AEAD, and Oribatida
- Institute of Applied Information Processing and Communications, TU Graz, Austria (2)

  Ascon and ISAP
- Submissions from the respective candidate teams (8):

  Gimli, KNOT, Romulus, Spook, Subterranean 2.0, WAGE, TinyJAMBU, and Xoodyak
- Submissions from other groups and independent researchers (2):

  Gimli, DryGASCON

Design Variants

Different variants corresponds to

- different algorithms of the same family described in a single submission to the NIST LWC standardization process
- different parameter sets, such as sizes of keys, nonces, tags, etc.
- support for AEAD vs. AEAD+Hash
- different hardware architectures, e.g., basic iterative, folded, unrolled, pipelined, etc.

72 variants

Minimum: 1, Maximum: 12, Average: 2.7

per hardware design submission

Benchmarking Flow – Part 1

1. Functional verification using GMU Known Answer Tests (KATs) not known to the designers in advance

2. Timing measurements

   a. 16 bytes, 64 bytes, 1536 bytes, N input blocks, N + d input blocks, with N = 4 and d = 1 or 2

   b. AD: AD only

      PT: Plaintext/Ciphertext only

      AD+PT: equal-size AD and Plaintext/Ciphertext

Benchmarking Flow – Part 2

3. Synthesis, Implementation, and Optimization of Tool Options

Xilinx:

Xilinx Vivado 2020.1 (lin64)

Minerva

Intel:

Intel Quartus Prime Lite Edition Design Software, ver. 20.1

ATHENa

Lattice Semiconductor:

Lattice Diamond Software v3.11 SP2

Synthesis: Lattice Synthesis Engine (LSE) or Synplify Pro

Xeda (new)

All results reported after placing & routing

Benchmarking Flow – Part 3

4. Calculation of Throughputs

5. Generation of graphs and tables

6. Analysis of results

Throughput vs. Area for Long Inputs
Artix-7 FPGA: Auth Encryption, Plaintext only

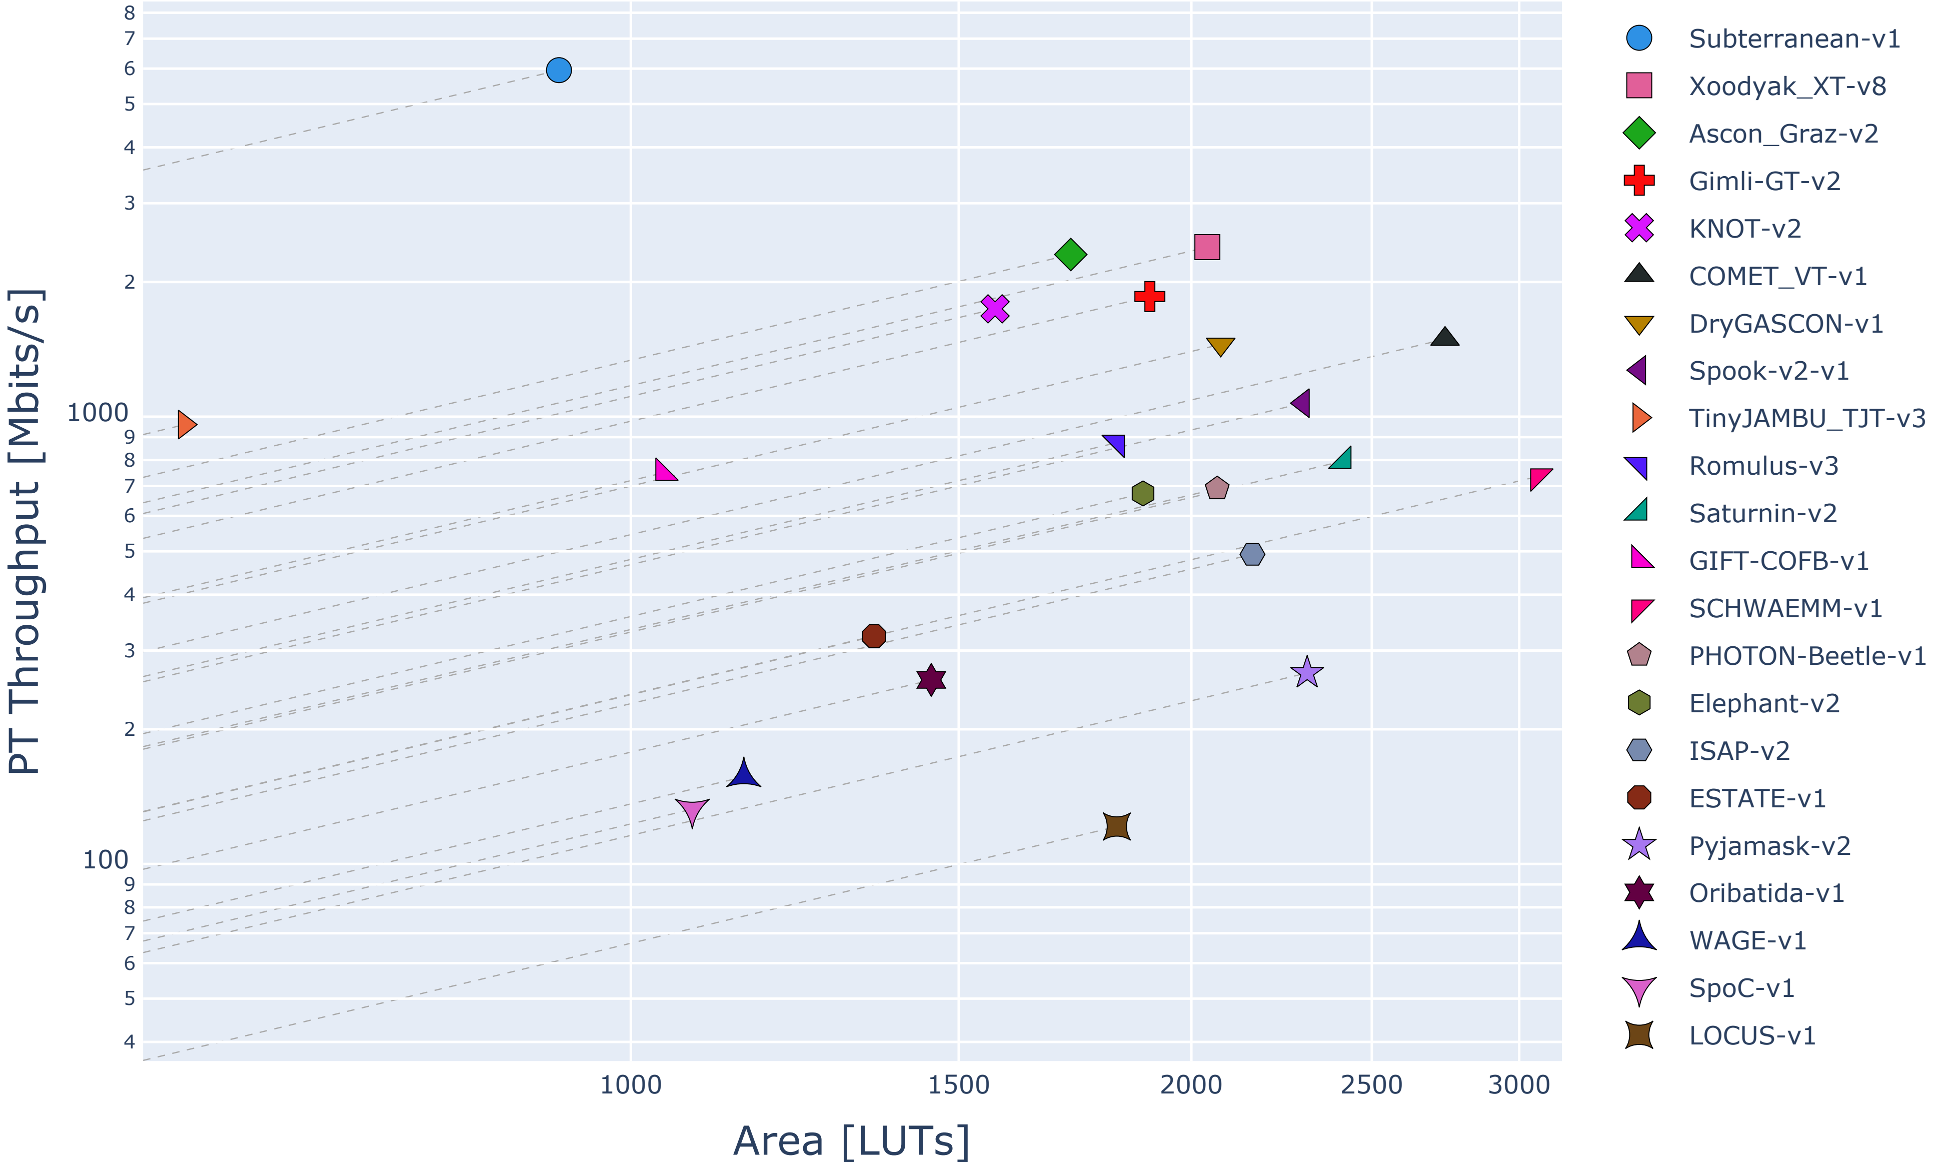

Throughput vs. Area for Long Inputs Artix-7 FPGA: Auth Encryption, AD only

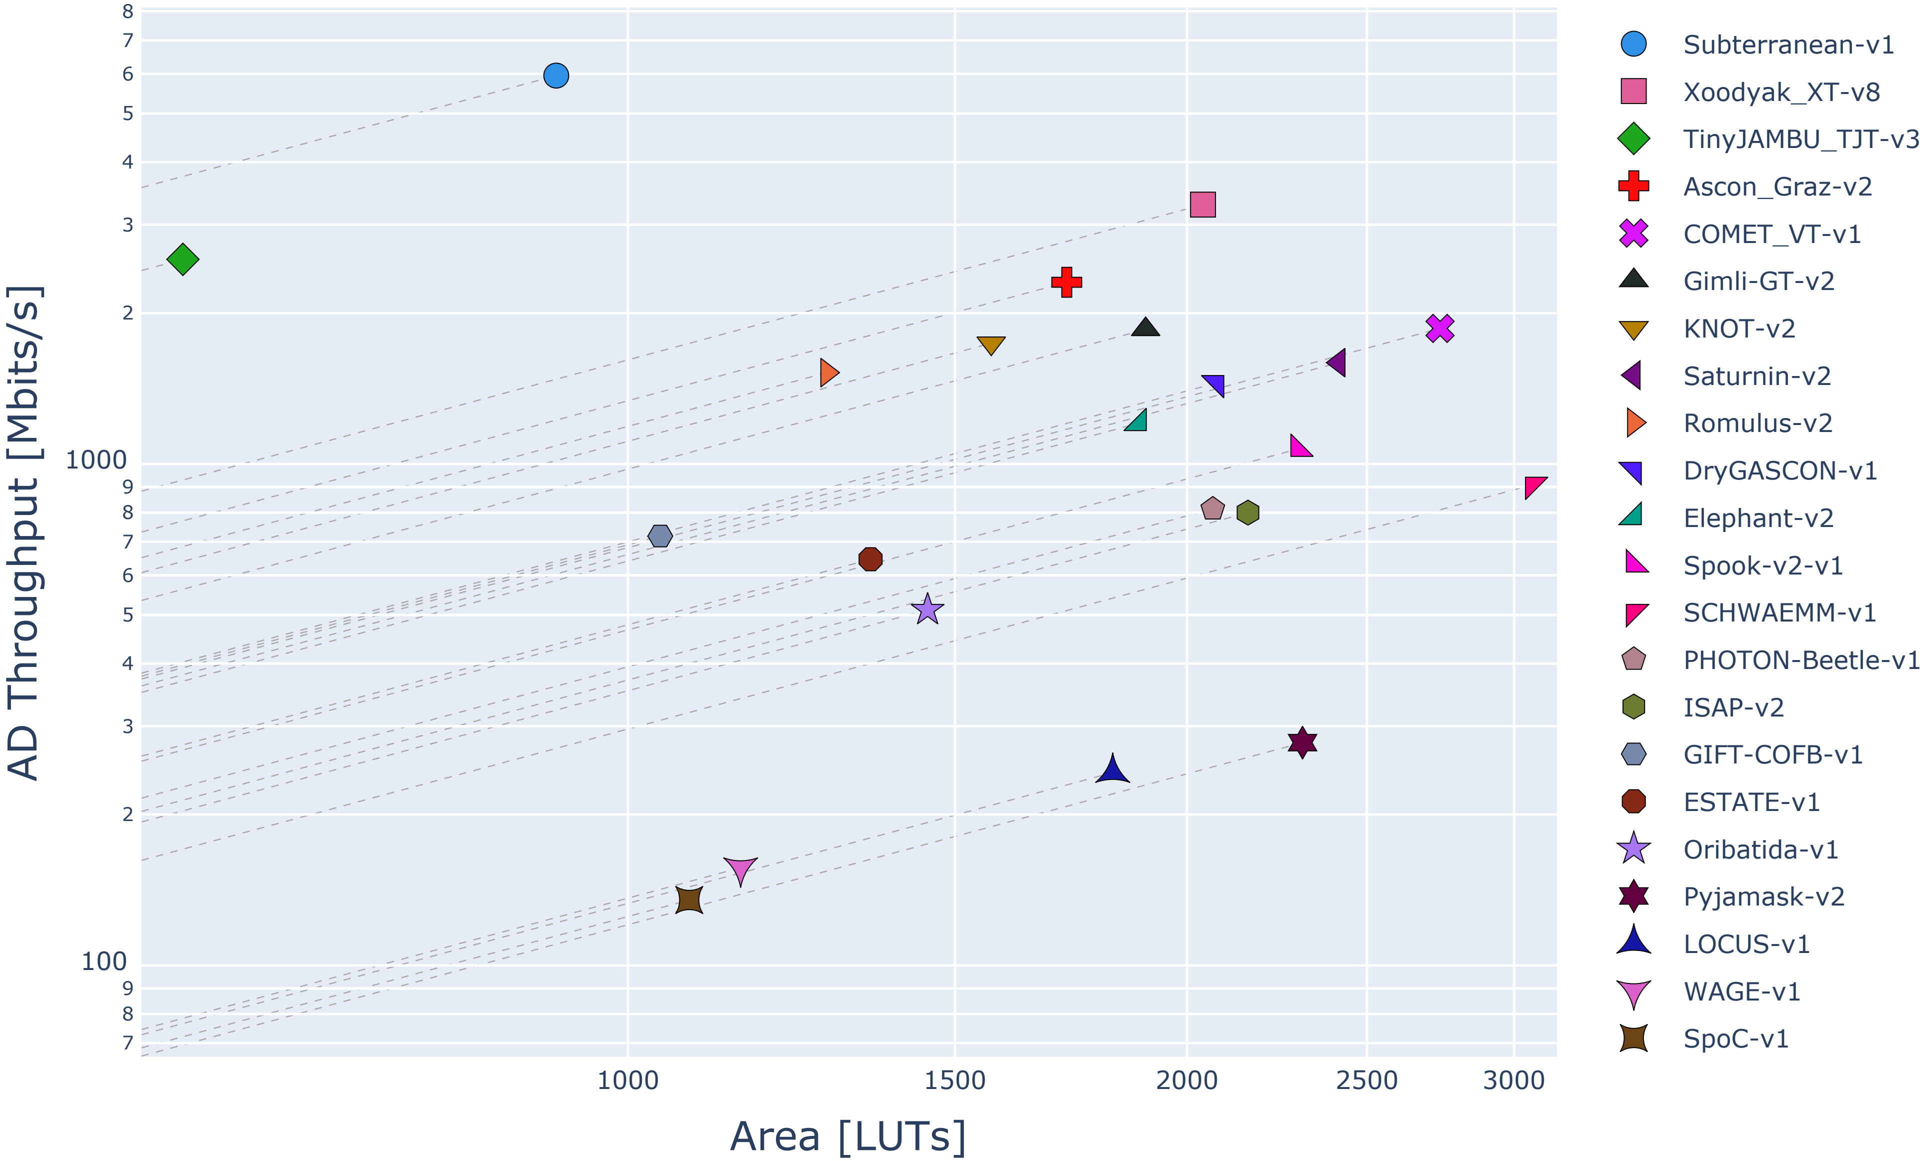

Throughput vs. Area for Long Inputs
Artix-7 FPGA: Hashing

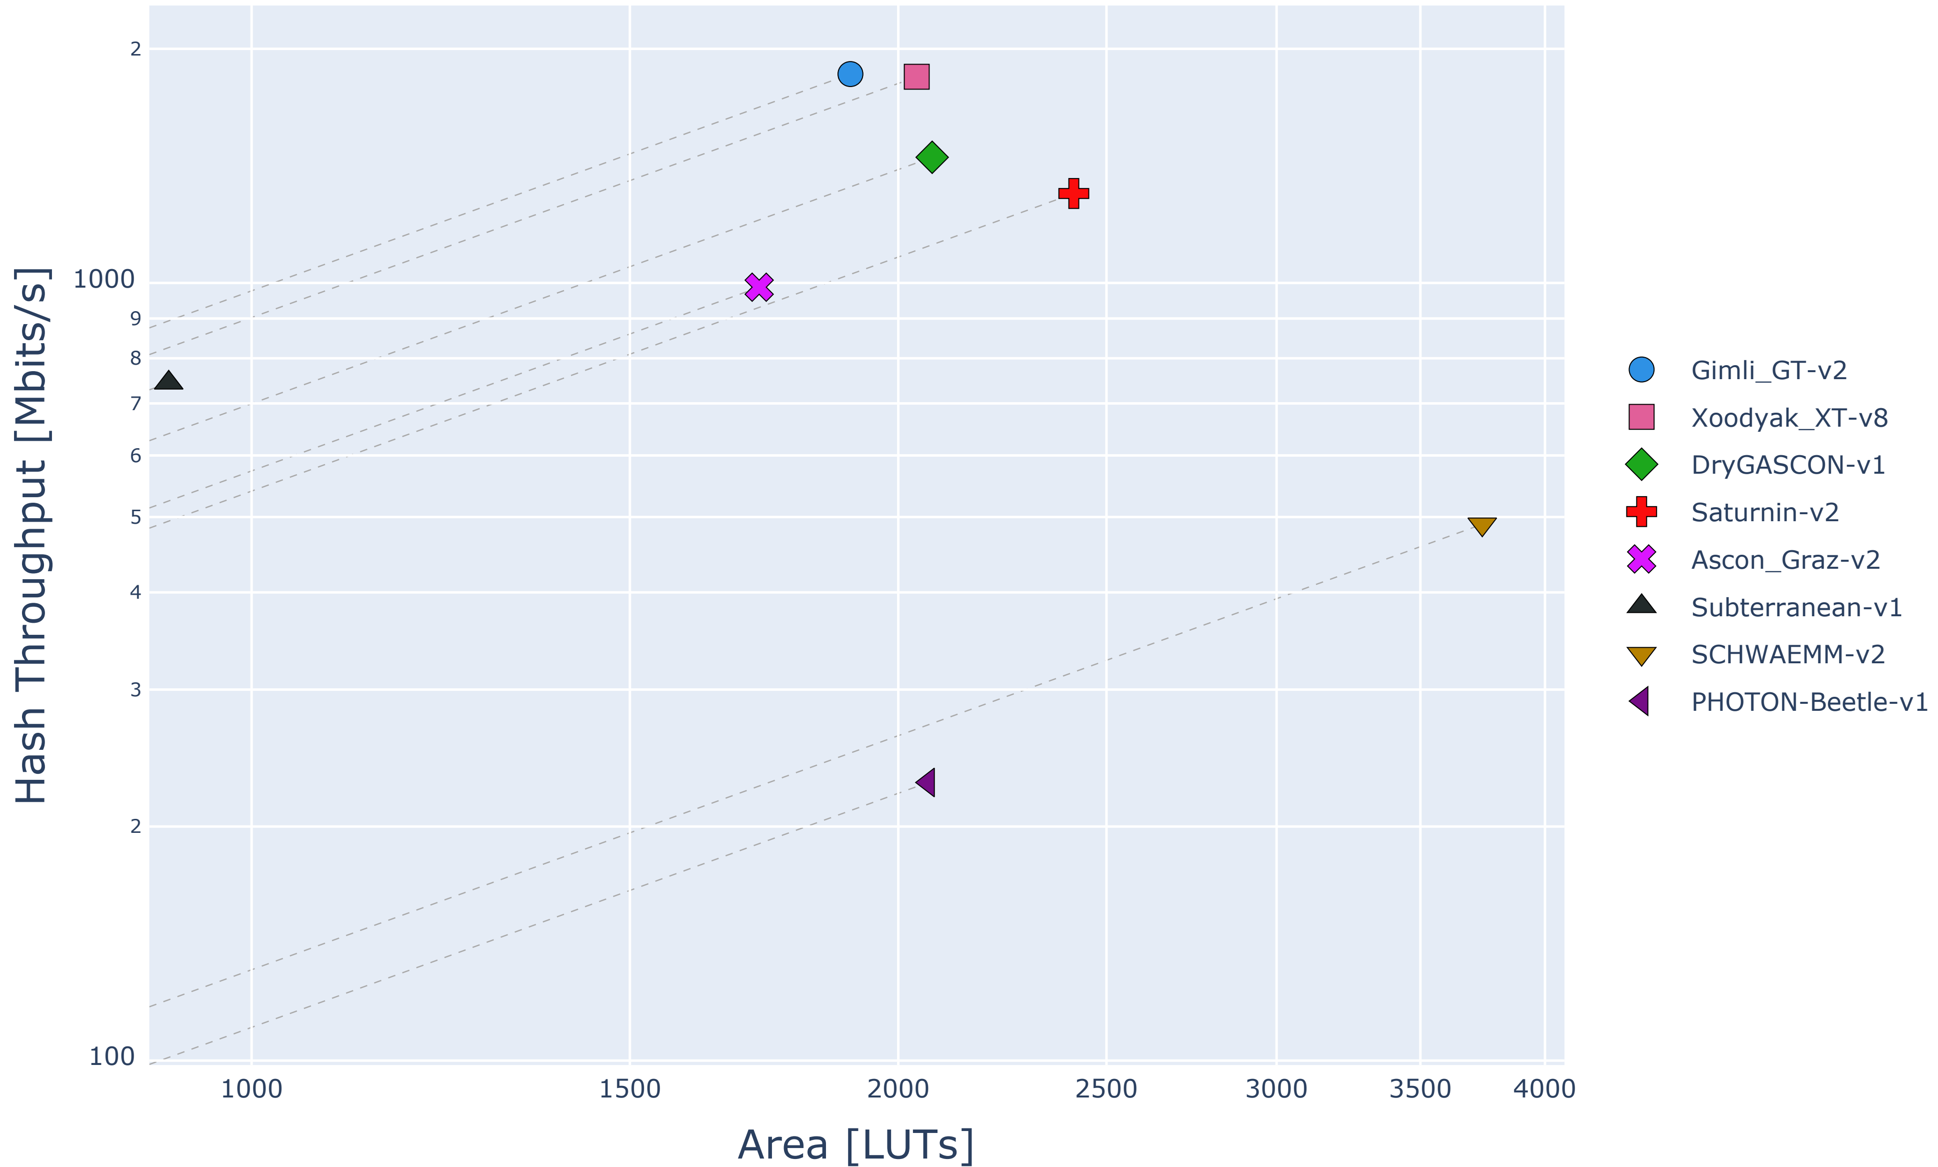

Design Space Exploration: Ascon
Artix-7 FPGA: Auth Encryption, Plaintext only

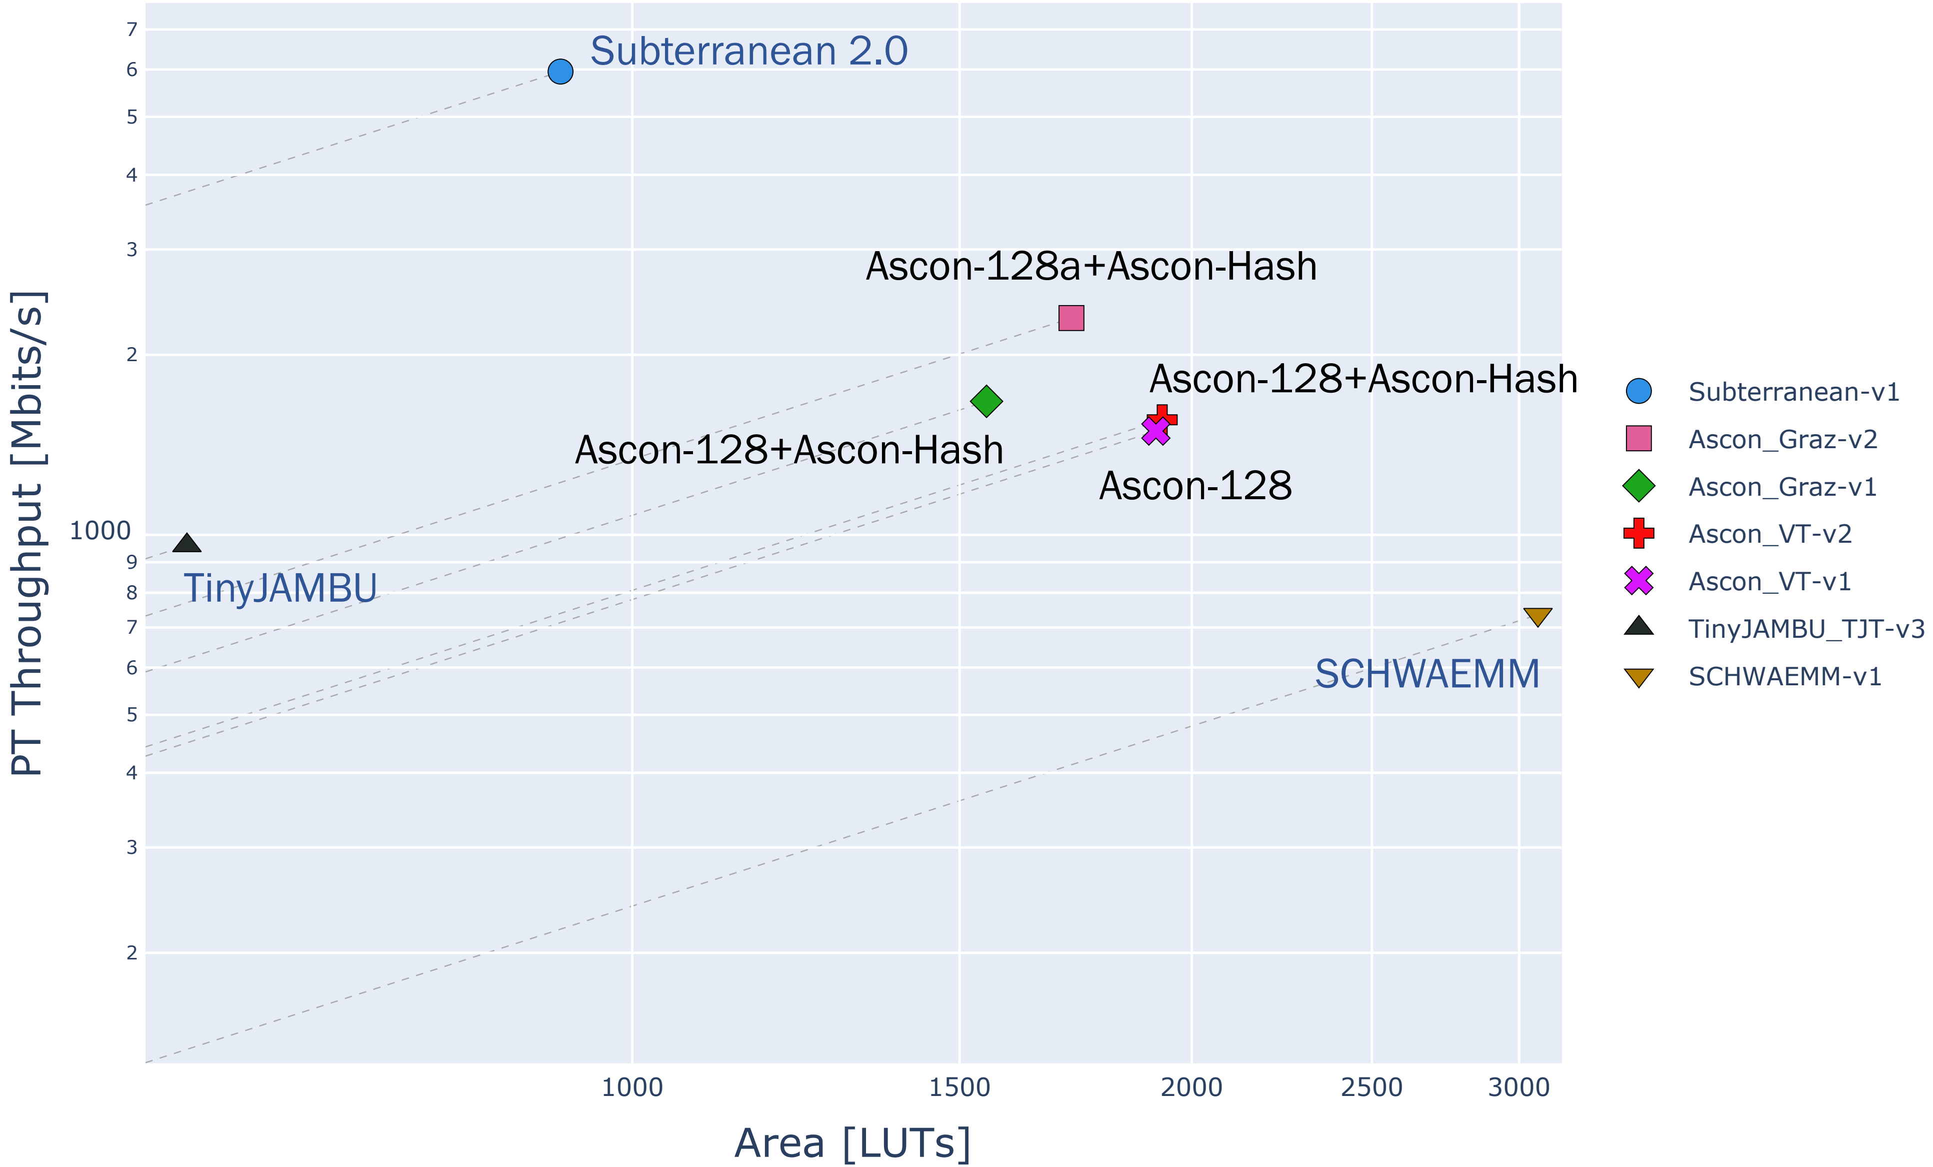

Design Space Exploration: KNOT
Artix-7 FPGA: Auth Encryption, Plaintext only

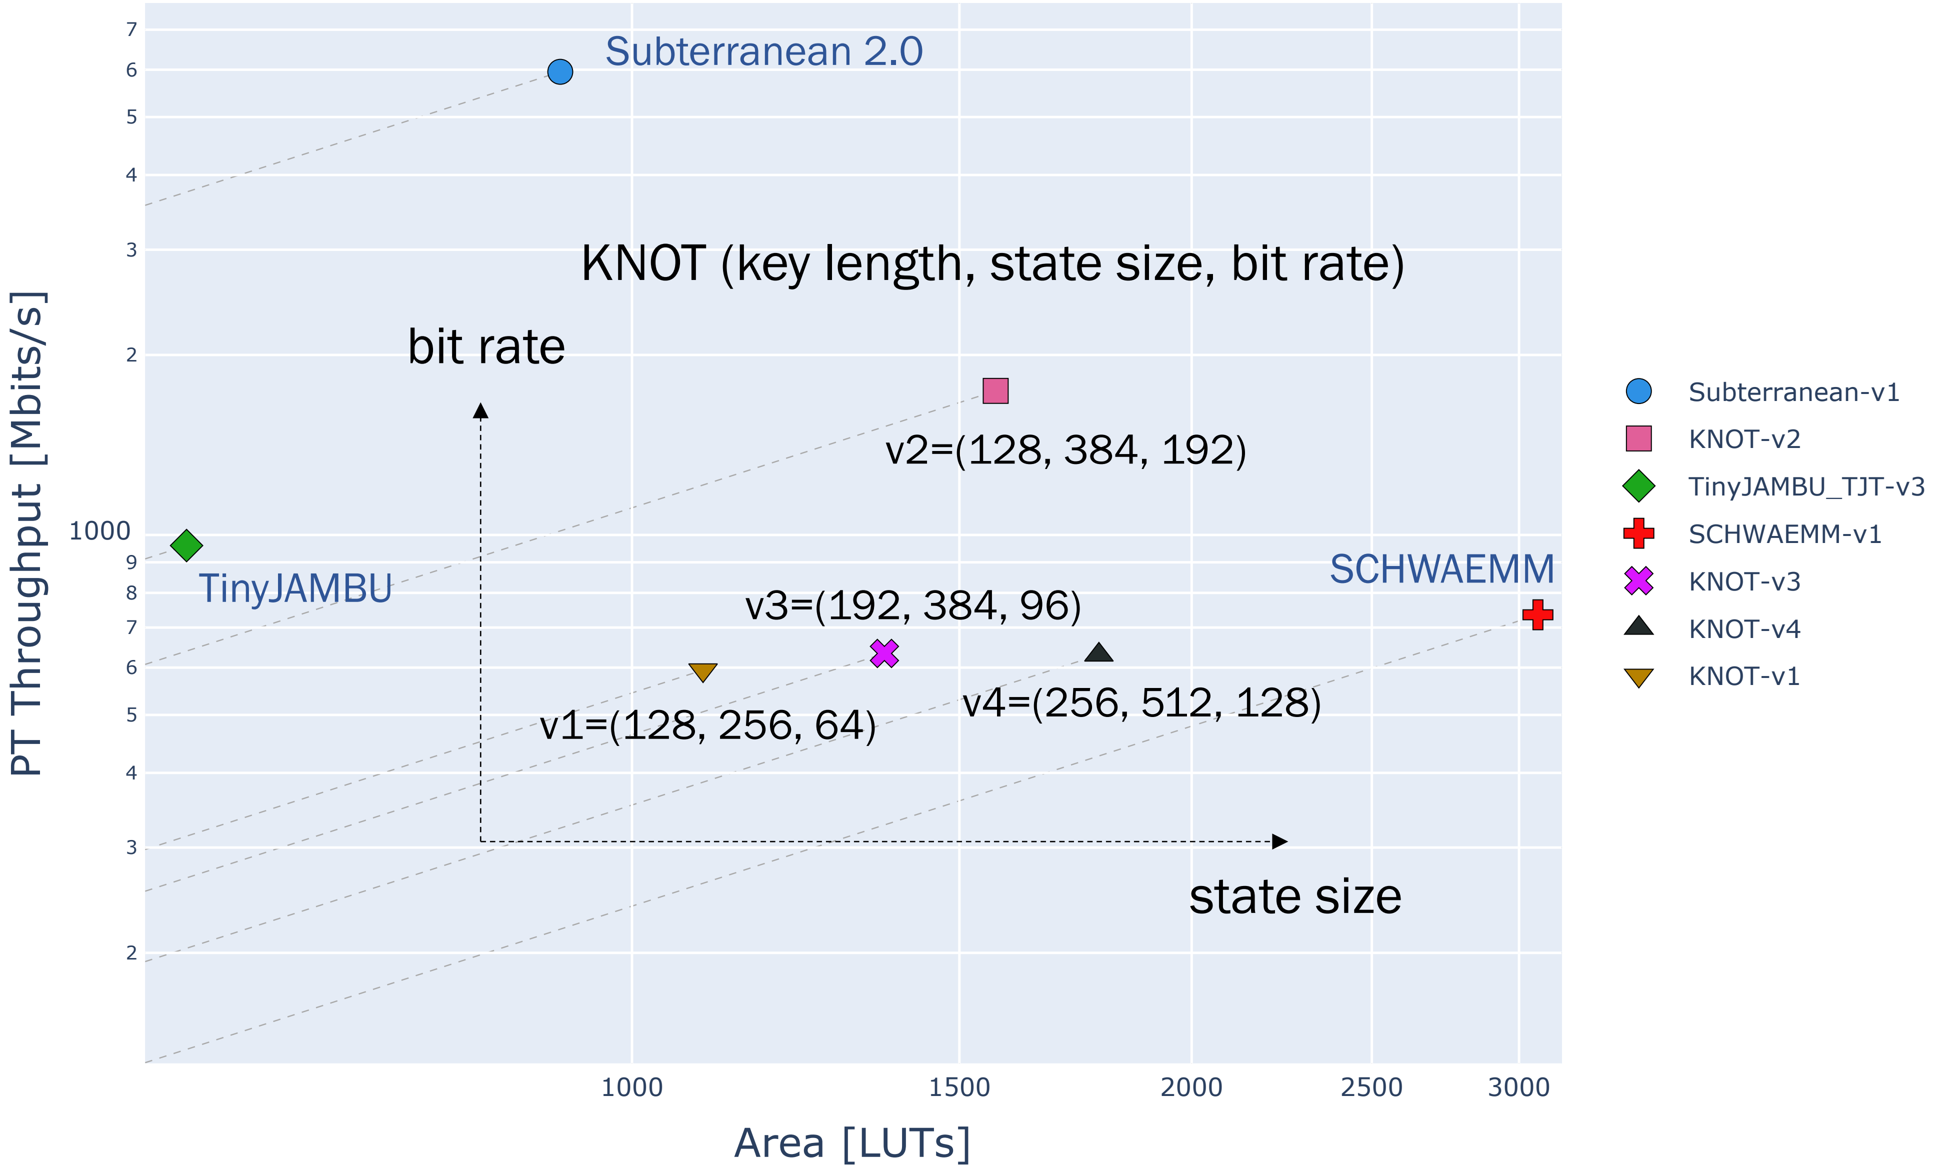

Design Space Exploration: Romulus
Artix-7 FPGA: Auth Encryption, Plaintext only

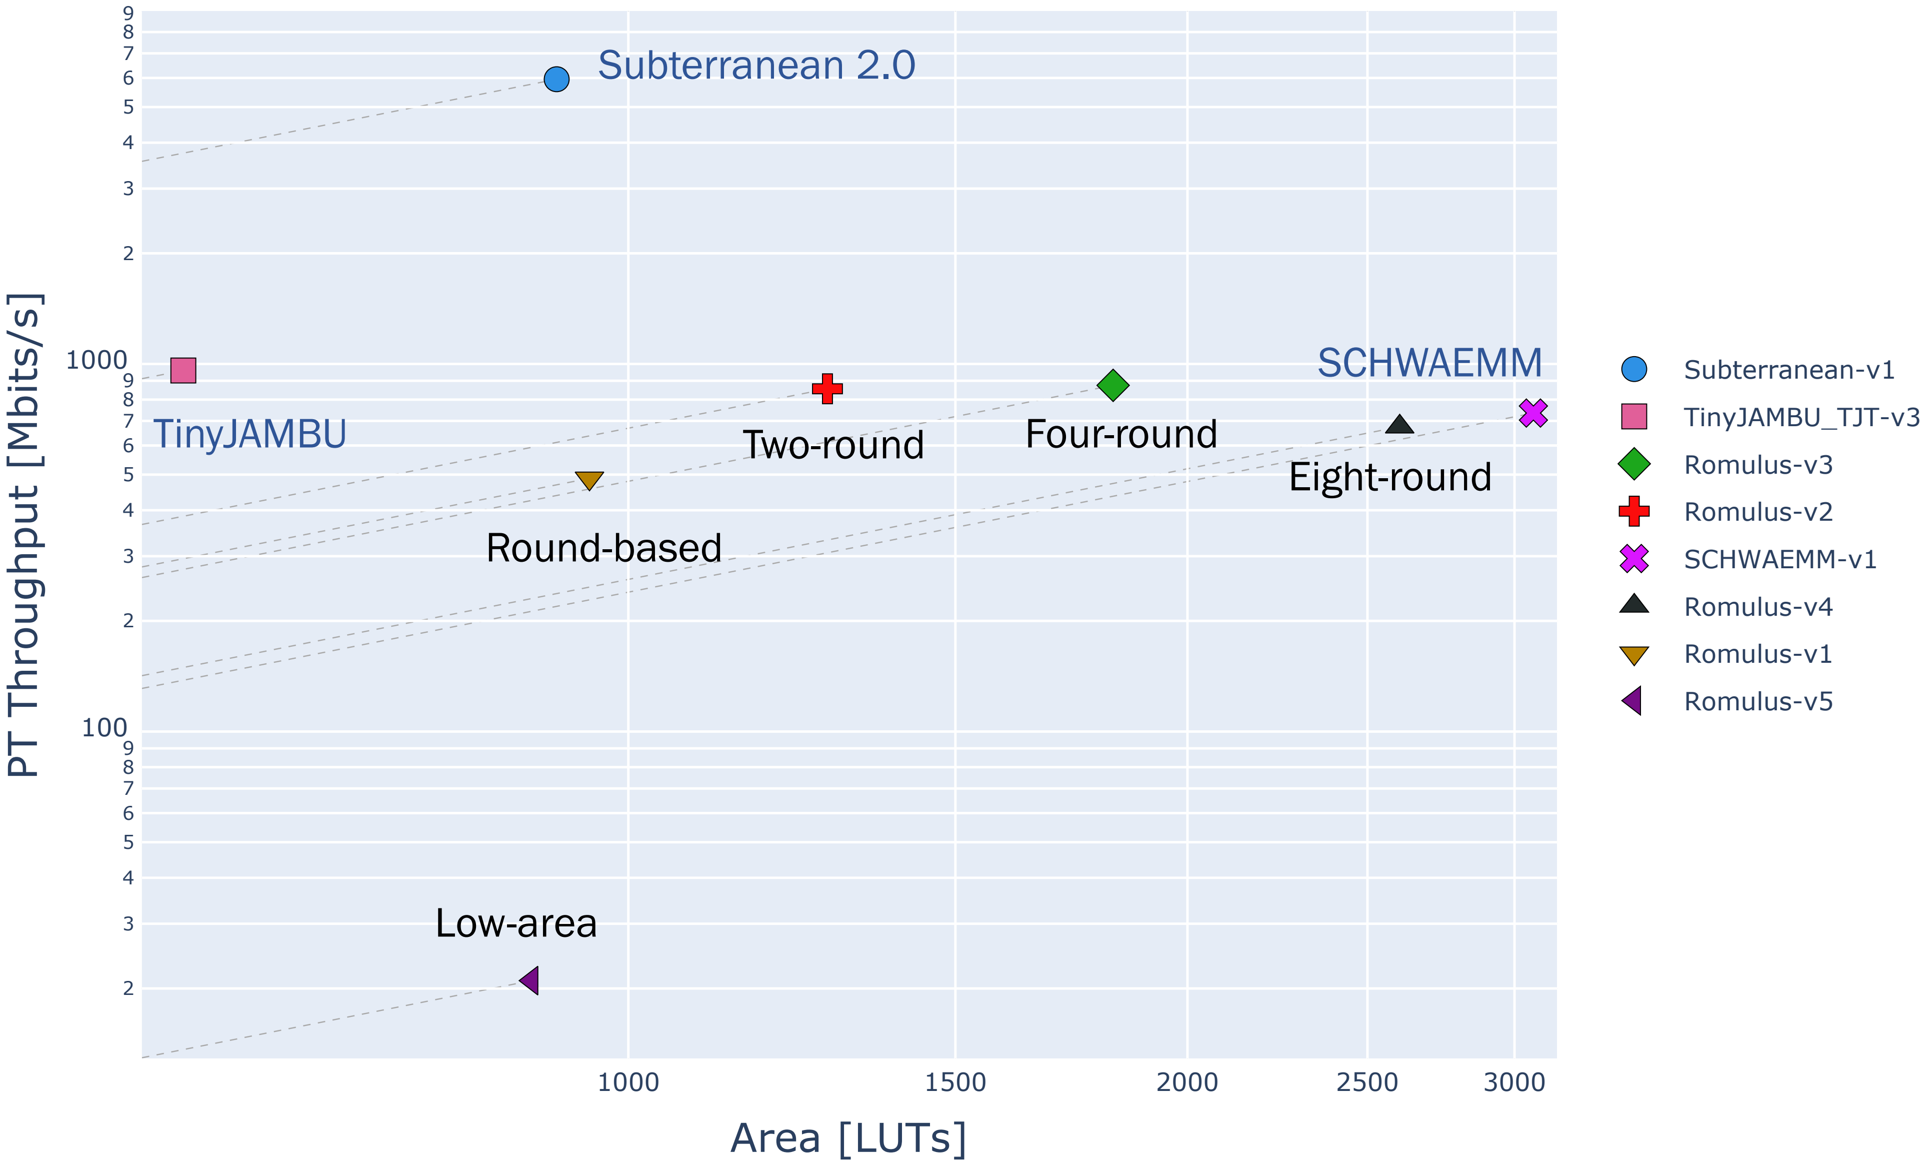

Dependence of Rankings on the Input Size
Artix-7 FPGA: Auth Encryption, Plaintext only

Position	Long	1536 B	64 B	16 B
1	Subterranean-v1	Subterranean-v1	Subterranean-v1	COMET_VT-v1
2	Xoodyak_XT-v8	Xoodyak_XT-v8	Ascon_Graz-v2	Subterranean-v1
3	Ascon_Graz-v2	Ascon_Graz-v2	Xoodyak_XT-v8	Ascon_Graz-v2
4	Gimli_GT-v2	Gimli_GT-v2	COMET_VT-v1	DryGASCON-v1
5	KNOT-v2	KNOT-v2	DryGASCON-v1	Xoodyak_XT-v8
6	COMET_VT-v1	COMET_VT-v1	TinyJAMBU_TJT-v3	TinyJAMBU_TJT-v3
7	DryGASCON-v1	DryGASCON-v1	KNOT-v2	Romulus-v2
8	Spook-v2-v1	Spook-v2-v1	Gimli_GT-v2	PHOTON-Beetle-v1
9	TinyJAMBU_TJT-v3	TinyJAMBU_TJT-v3	Romulus-v2	KNOT-v2
10	Romulus-v2	Romulus-v2	Spook-v2-v1	Gimli_GT-v2
11	Saturnin-v2	Saturnin-v2	PHOTON-Beetle-v1	Elephant-v2

Color code: Higher position for smaller messages Lower position for smaller messages

Dependence of Rankings on the Input Size
Artix-7 FPGA: Auth Encryption, Plaintext only

Position	Long	1536 B	64 B	16 B
12	GIFT-COFB-v1	GIFT-COFB-v1	GIFT-COFB-v1	GIFT-COFB-v1
13	SCHWAEMM-v1	SCHWAEMM-v1	Elephant-v2	ESTATE-v1
14	PHOTON-Beetle-v1	PHOTON-Beetle-v1	SCHWAEMM-v1	Spook-v2-v1
15	Elephant-v2	Elephant-v2	Saturnin-v2	SCHWAEMM-v1
16	ISAP-v2	ISAP-v2	ESTATE-v1	Oribatida-v1
17	ESTATE-v1	ESTATE-v1	Oribatida-v1	Saturnin-v2
18	Pyjamask-v2	Pyjamask-v2	ISAP-v2	LOCUS-v1
19	Oribatida-v1	Oribatida-v1	Pyjamask-v2	SpoC-v1
20	WAGE-v1	WAGE-v1	SpoC-v1	ISAP-v2
21	SpoC-v1	SpoC-v1	LOCUS-v1	Pyjamask-v2
22	LOCUS-v1	LOCUS-v1	WAGE-v1	WAGE-v1

Color code: Higher position for smaller messages   Lower position for smaller messages

Tentative Conclusions

Highest Throughput for #LUTs < 2500

Plaintext Only

1.	Subterranean 2.0	~6 Gbit/s
2.	Xoodyak	2-3 Gbit/s
3.	Ascon	
4.	Gimli	1-2 Gbit/s
5.	KNOT	
6.	DryGASCON	
7.	Spook	

AD Only

1.	Subterranean 2.0	~6 Gbit/s
2.	Xoodyak	3-4 Gbit/s
3.	Ascon	2-3 Gbit/s
4.	**TinyJAMBU**	
5.	Gimli	1-2 Gbit/s
6.	KNOT	
7.	**Saturnin**	
8.	**Romulus**	
9.	DryGASCON	
10.	**Elephant**	
11.	Spook	

Future Work in Round 2

Phase 3:

Nov. 9, 2020: 3rd submission deadline
Nov. 30, 2020: Final version of the report

- Hopefully 85%-100% of candidates!
- Improved designs
- More design space explorations
- A new tool supporting the derivation of formulas for the execution times
- Evaluation in terms of Power consumption and Energy per bit

Living Report from Round 2

https://cryptography.gmu.edu/athena

Cryptology ePrint Archive: Report 2020/1207

~90 pages, ~25 graphs, ~60 tables

Released on GMU website: Sep. 26, 2020
Posted on ePrint: Oct. 2, 2020
Phase 2 Report available tomorrow, October 21, 2020
Regular updates
Changelog at the end of the document

Q & A