

KU LEUVEN





Hardware: an essential partner to cryptography

IACR Distinguished Lecture, **UPDATE**
 Prof. Dr. Ir. Ingrid Verbauwhede
 KU Leuven, COSIC

Eurocrypt 2022 – June 1 2022
 Croatia – June 8, 2023

Slide acknowledgement: All past and present PhD students!





1

COSIC = Research group










Claudia Diaz	Nele Mentens	Kris Myny	Bart Preneel	Vincent Rijmen	Nigel Smart	Ingrid Verbauwhede	Fre Vercauteren
Privacy Technologies	FPGA Hardware Acceleration	Beyond CMOS technologies	Cryptography Applied Security	Symmetric Key	Public Key Protocols	Crypto HW Security	Algebra Crypto



2

Join us



3

KU LEUVEN

3



KU LEUVEN

4

COSIC: “an” ENCRYPTION POWERHOUSE RISES

- Wall Street Journal dd. 10/12/2015: “In Belgium, an Encryption Powerhouse Rises, University of Leuven has become a battleground in the fight between privacy and surveillance”
- “Packed with hardware and laptop-wielding students in jeans and sneakers, COSIC’s labs develop new encryption for corporate clients, or test their in-house antihacking technology.”



<https://www.wsj.com/articles/in-belgium-an-encryption-powerhouse-rises-1449791014>

KU LEUVEN

5

Outline

- Position of cryptography in the design of embedded systems
 - Root of trust & secure composition
- Cryptography relies on hardware because it needs:
 - Performance (see DES chip)
 - Secure implementation: protection against side-channel, fault attacks
 - Secure key storage (PUFs)
 - Quality random number generators
 - Acceleration of new crypto: COED and FHE
- Challenges for crypto to work on
- Conclusions

6

KU LEUVEN

6

NEXT GENERATION EMBEDDED SYSTEMS

7

KU LEUVEN

7

Automotive

“Networked embedded systems interacting with the environment”

WELKE SOORT AUTOCHIPS MAAKT MELEXIS?

- Jandrijging
- Chassis, carrosserie en veiligheid
- Pcil ruftekwissenloestof
- Rufftwasser motor
- Audi-condens voornut
- Schakelaar motorapodot
- Stuurberediging
- Luchtgasrembrude
- Parkapodotremmer
- Veroneringbak
- Elektrische handrem
- Motorcooling ventitor
- Ingebouwde opodot
- Koppelingsschakelaar
- Schakelapodotremmer
- Remschakelaar
- Detectie remvoelstodpel
- Regen-lichtsensor
- Schuifmotor
- Energie-efficiënte klimaatcontrole
- Skutelloze toegang
- Schakelaar voor deurapodot
- Motor voor rufte
- Durfremdele
- Koffierapodot
- Motor openen koffier
- Stoflampe
- Richtigingapodot
- Verlichting
- Batterij-monitoring systeem
- Batterij temp. management
- Spanningapodot
- Elektrische handrem
- Sensor veiligheidstriem
- Stoelapodotremmer
- Verkele stoelen
- Stoelapodotremmer

ANDY GREENBERG SECURITY 07.21.15 06:00 AM

HACKERS REMOTELY KILL A JEEP ON THE HIGHWAY—WITH MELEXIS

- Networked → Secure, authenticated communication, low latency
- Embedded → compact (no external memory), cheap, no batch processing
- Interacting with environment
 - LOW latency
 - Compact
- Resistant to attacks

8

KU LEUVEN

[De Tijd, February 2, 2022]

8


How to evaluate security? Where to start?



Tesla Model X key fob (2020)
<https://youtu.be/clrNuBb3myE>

Tesla Model S key fob (2018)
<https://youtu.be/aVIYuPzmJoY>

[Lennert Wouters, COSIC]



Passive Keyless Entry and Start System:

- Wireless challenge response system
- **No Mutual authentication (model S)**
- **Weak crypto (model S)**
- **Secure element, but problems with protocol (model X)**
- Off the shelf radios and components

9

KU LEUVEN

9

TRUST AND TRUST BOUNDARIES

10

KU LEUVEN

10

Trust Definition

Trust (R. Anderson in “Security Engineering”, after NSA):

- “Trusted system or component is one whose failure can break the security policy, while a trustworthy system or component is one that won’t fail.”

Trust (Trusted Computing Group):

- “An entity can be trusted if it always behaves in the expected manner for the intended purpose.”

Loosely stated: if trusted system or component fails, then bad things can happen.

Goal of security: **minimize** what needs to be trusted

How does cryptography fit in this context?

11

KU LEUVEN

11

What is the root of trust?

- For network system: router box
- For secure boot: the TPM or SE
- For OS designer: the architecture/micro-architecture of a processor
- For cryptographer: the VHDL or Verilog code on FPGA
- For IOT devices: attack resistance (side-channel, fault, manipulation, etc.)
- For digital designer: the standard cells or the technology

KU LEUVEN

12

HOW: DESIGN METHOD

DECOMPOSE IN COMPONENTS

- Application: secure communication
- Cryptography: public key, secret key, post-quantum,
- Architecture: Hardware/Software platform, Sancus
- Micro-architecture: crypto co-processors, instruction set extension,
- Logic circuits and (secure) memory
- TRNGs and PUFs
- Technology

[P. Schaumont, I. Verbauwhede, "Design methods for security and Trust, DATE2007]

KU LEUVEN

13

Recent:

- US CHIPS and Science Act
- September 2022

★★★

REPORT TO THE PRESIDENT
Revitalizing the U.S.
Semiconductor Ecosystem

Executive Office of the President
President's Council of Advisors on
Science and Technology

September 2022

LEUVEN

14

(e) Semiconductors and System Security

Criminal and state-sponsored cyber-attacks pose increasing threats to the United States. To enable the implementation of secure systems, every aspect of the system must be considered including sensors, data converters, computing, memory storage, and communications, while providing robustness against side-channel attacks and ensuring security of supply chains. There is a tremendous opportunity for the design of secure semiconductor chips. To maximize effectiveness, security must be pursued as an integral part of design, not as an add-on after the chip is designed.

Academia, industry and government stakeholders have an opportunity to standardize a trusted approach for systems implementation. The specific opportunity is to bring together algorithm and software/systems designers with chip designers in a center of excellence, to develop the next generation of secure systems. Although open-source security approaches are the best for innovation and transparency, they remain unpalatable for the industry. We must address this reluctance in a way that enables the United States to continue to be the global leader in standardized security approaches.

We envision a research agenda in this area that should include, but is not limited to, the following: (1) design for fully secure end-to-end hardware and software solutions that are secure against various forms of attacks on operation, data, and communications; (2) security in the chip design tool chain that would enable end-to-end security solutions to be verified by design; (3) secure hardware supply chain covering chip fabrication, packaging, and system integration; (4) implementation of post-quantum cryptography; (5) implementation of low-power cryptography for secure communications and transactions; and (6) other privacy preserving hardware implementations for processing encrypted data.

VEN

15

We envision a research agenda in this area that should include, but is not limited to, the following: (1) design for fully secure end-to-end hardware and software solutions that are secure against various forms of attacks on operation, data, and communications; (2) security in the chip design tool chain that would enable end-to-end security solutions to be verified by design; (3) secure hardware supply chain covering chip fabrication, packaging, and system integration; (4) implementation of post-quantum cryptography; (5) implementation of low-power cryptography for secure communications and transactions; and (6) other privacy preserving hardware implementations for processing encrypted data.

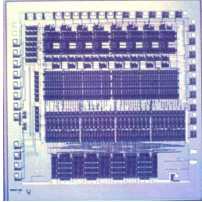
16

KU LEUVEN

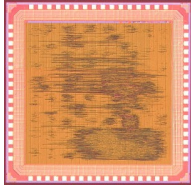
16

DES, AES, ECC, SABER dedicated ASICs

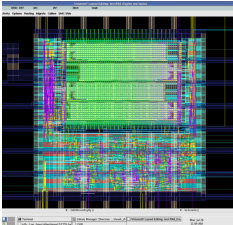
- Feasibility:** what is feasible, throughput, latency, power (cooling), energy (battery lifetime) etc.



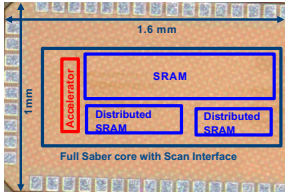
DES



Rijndael



ECC



Saber

- Next: light weight crypto, COED, FHE, ...

KU LEUVEN

17

Wide range of design options!

HW		HW-SW			SW		
ASIC	FPGA	Domain specific	CO-proc	DSP	VLIW GPU	CPU	
Area efficiency							
High							Low
Performance/Energy unit							
Low							High
Programmability							

Energy – throughput– cost - flexibility trade-off
Almost always: HW-SW co-design!

KU LEUVEN

18

Throughput – Energy numbers

AES 128bit key 128bit data	Throughput	Power	Figure of Merit (Gb/s/W = Gb/J)
0.18um CMOS	3.84 Gbits/sec	350 mW	11 (1/1)
FPGA [1]	1.32 Gbit/sec	490 mW	2.7 (1/4)
Intel ISA for AES [6]	32 Gbit/sec	95 W	0.34 (1/33)
ASM StrongARM [2]	31 Mbit/sec	240 mW	0.13 (1/85)
Asm Pentium III [3]	648 Mbits/sec	41.4 W	0.015 (1/800)
C Emb. Sparc [4]	133 Kbits/sec	120 mW	0.0011 (1/10.000)
Java [5] Emb. Sparc	450 bits/sec	120 mW	0.0000037 (1/3.000.000)

[1] Amphion CS5230 on Virtex2 + Xilinx Virtex2 Power Estimator
 [2] Dag Ane Osvik: 544 cycles AES – ECB on StrongArm SA-1110
 [3] Helger Lipmaa PIII assembly handcoded + Intel Pentium III (1.13 GHz) Datasheet
 [4] gcc, 1 mW/MHz @ 120 Mhz Sparc – assumes 0.25 u CMOS
 [5] Java on KVM (Sun J2ME, non-JIT) on 1 mW/MHz @ 120 Mhz Sparc – assumes 0.25 u CMOS
 [6] Shay Gueron, Intel

[P. Schaumont, and I. Verbauwhede, "Domain specific codesign for embedded security," Computer 36(4), pp. 68-74, 2003.]

KU LEUVEN

19


We envision a research agenda in this area that should include, but is not limited to, the following: (1) design for fully secure end-to-end hardware and software solutions that are secure against various forms of attacks on operation, data, and communications; (2) security in the chip design tool chain that would enable end-to-end security solutions to be verified by design; (3) secure hardware supply chain covering chip fabrication, packaging, and system integration; (4) implementation of post-quantum cryptography; (5) implementation of low-power cryptography for secure communications and transactions; and (6) other privacy preserving hardware implementations for processing encrypted data.

KU LEUVEN

20

Side-channel and fault attacks

- Many types of side-channel analysis
 - Power, Electro Magnetic (EM), Time,
 - Micro-architectural side-channel: cache, transient execution attacks
- Many types of fault or active attacks:
 - EM, laser, clock, voltage glitch, etc.
- Local or remote
- Combined attacks



21

21


Measurement methods

Contact power measurements:

- shunt resistors
- current probes

Cost: 150- 5000€

Freq: kHz – MHz range




Contactless power measurements:

- EM probes

Cost: 2000 - 25000€

Freq: kHz – GHz range




[picture credit: Langer]

EM measurements:


- EM probes

Cost: 2000 - 25000€

Freq: kHz – GHz range



[picture credit: Langer]



22

Research challenges for cryptography

- Goal: introduce new research topics, improve existing ones
- **Challenge 1: masking is hard in practice**
- Challenge 2: masking is expensive
- Challenge 3: Possibilities of PUFs
- Challenge 4: Random number generation
- Challenge 5: NEW – Fully Homomorphic Encryption

23

KU LEUVEN

23

Countermeasure: masking

- Types of masking
 - Boolean
 - Arithmetic
 - Inner product
 - Threshold
 - ...
 - Two experiments:
 - Symmetric key: AES masking on micro controllers
 - Public key: Post-quantum masking of lattice based encryption
- All start from similar leakage MODEL:
Shares leak independently**
- All require randomness**

24

KU LEUVEN

24

Masking in practice is HARD

- Experiment: **first** order SW masked AES evaluated for:
 - Side-channel leakage
 - Timing
 - Randomness requirements

Paper title	Published venue	masking method
Provably Secure Higher-Order Masking of AES	CHES 2010	boolean
Higher order masking of look-up tables	Eurocrypt 2014	boolean
All the AES You Need on Cortex-M3 and M4	SAC 2016	boolean
Consolidating Inner Product Masking	Asiacrypt 2017	inner product
First-Order Masking with Only Two Random Bits	CCS-TIS 2019	boolean
Side-channel Masking with Pseudo-Random Generator	Eurocrypt 2020	boolean
Detecting faults in inner product masking scheme	JCEN 2020	inner product
Fixslicing AES-like Ciphers	TCHES 2021	boolean

[A. Becker, L. Wouters, Cosade 2022]

25

KU LEUVEN

25

Results [Cosade 2022]

- Key recovery with first order attack ●
- Incorrect TRNG instantiations ●
- Benchmarking issues ●
- Software bugs ●

	Paper title	Published venue	masking method
●	Provably Secure Higher-Order Masking of AES	CHES 2010	boolean
●	Higher order masking of look-up tables	Eurocrypt 2014	boolean
● ●	All the AES You Need on Cortex-M3 and M4	SAC 2016	boolean
	Consolidating Inner Product Masking	Asiacrypt 2017	inner product
●	First-Order Masking with Only Two Random Bits	CCS-TIS 2019	boolean
● ● ●	Side-channel Masking with Pseudo-Random Generator	Eurocrypt 2020	boolean
	Detecting faults in inner product masking scheme	JCEN 2020	inner product
●	Fixslicing AES-like Ciphers	TCHES 2021	boolean

26

KU LEUVEN

26

Correlation Power Analysis (CPA)



- All implementations compiled using given makefile
- Only inserted triggers
- Textbook first order CPA:
 - SBOX in or output
 - Hamming Weight leakage, or single bit when bitsliced
 - 20k traces
- No claims about the mathematical concepts or proofs

27

KU LEUVEN

27

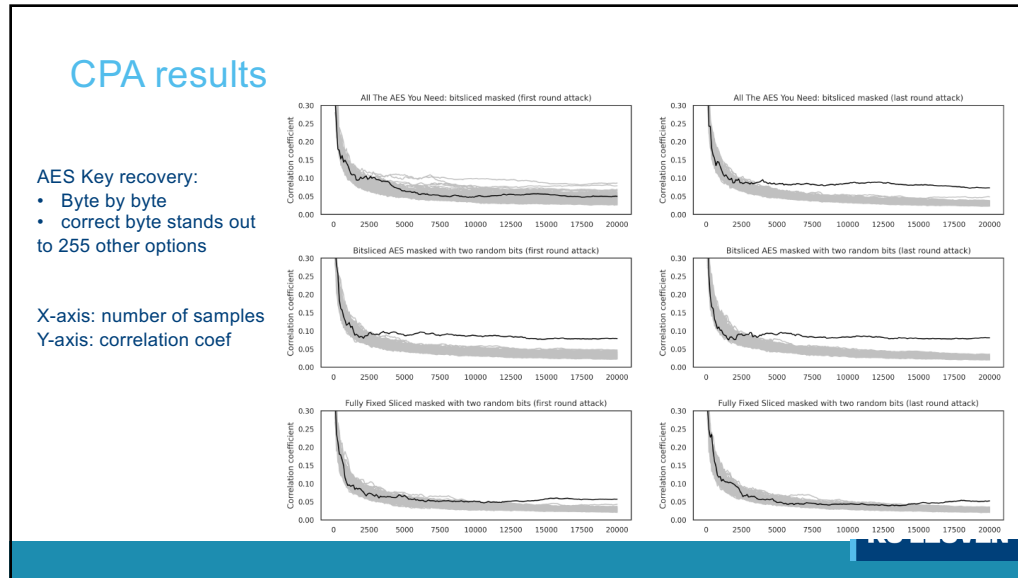
Set-up in the lab



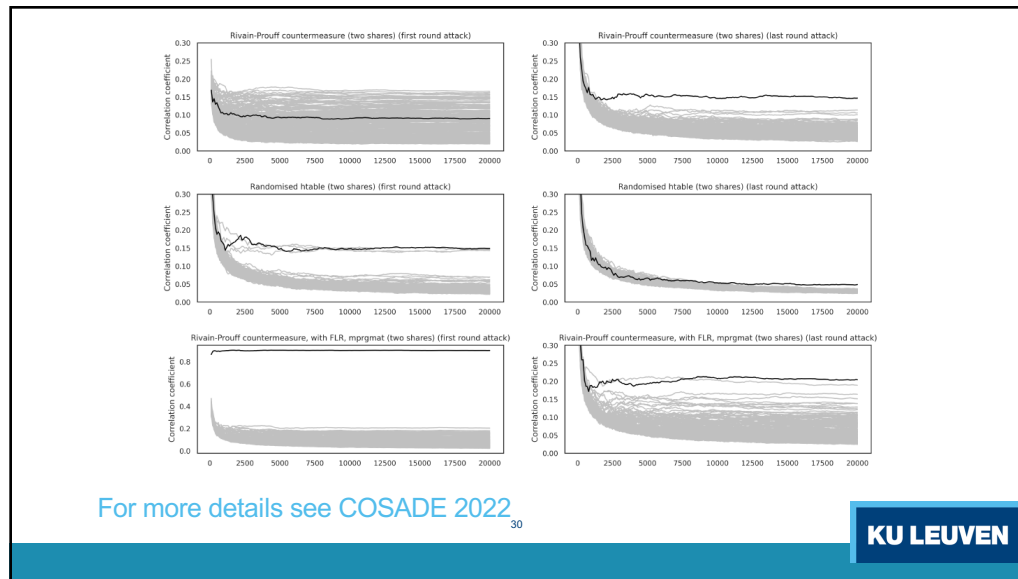
28

KU LEUVEN

28



29



30

Cause: violation of assumptions

- Assumption: shares leak independently
- Leakage caused by the microcontroller breaks this assumption
 - Assume share A is in r0
 - Move share B into r0 (and overwrite share A)
 - Information on $A \oplus B$ is leaked!
- Complex processors: transient execution
- Compiler optimizations
- Coupling through power and ground network
- Below 60nm CMOS 'static' leakage ³¹

**EDA message:
TOOLS could help here!**

KU LEUVEN

31

Research challenges for cryptography

- Goal: introduce new research topics, improve existing ones
- Challenge 1: masking is hard in practice
- **Challenge 2: masking is expensive**
- Challenge 3: Possibilities of PUFs
- Challenge 4: Random number generation
- Challenge 5: NEW – Fully Homomorphic Encryption

32

KU LEUVEN

32

We envision a research agenda in this area that should include, but is not limited to, the following: (1) design for fully secure end-to-end hardware and software solutions that are secure against various forms of attacks on operation, data, and communications; (2) security in the chip design tool chain that would enable end-to-end security solutions to be verified by design; (3) secure hardware supply chain covering chip fabrication, packaging, and system integration; (4) implementation of post-quantum cryptography; (5) implementation of low-power cryptography for secure communications and transactions; and (6) other privacy preserving hardware implementations for processing encrypted data.


33

KU LEUVEN

33

Lattice Based Post-quantum crypto (NIST)

- KEM = key generation, encapsulation, decapsulation
- CCA secure: Fujisaki – Okamoto transformation
- Similar for
 - Kyber
 - Saber



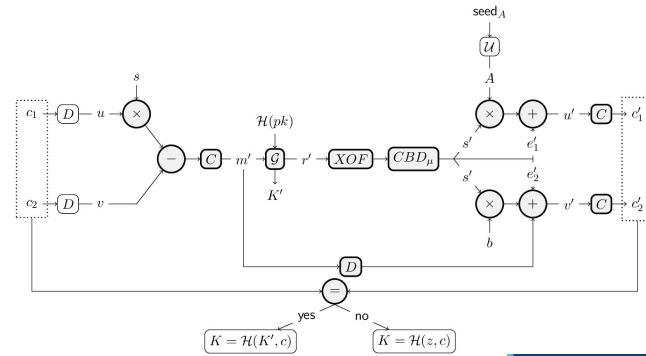
34

KU LEUVEN

34

Cost of decapsulation

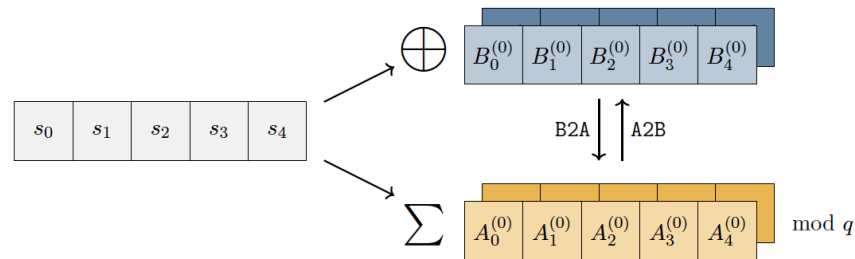
- Expensive parts: multiplication, hash, sampling
- Saber vs Kyber
 - Very similar
 - Power of two $q=2^{13}$ vs $q=3329$
 - MLWR vs MLWE implicit vs explicit error addition



35

35

Arithmetic and Boolean masking



Conversion is: Arithmetic to Boolean (A2B) or Boolean to Arithmetic (B2A)

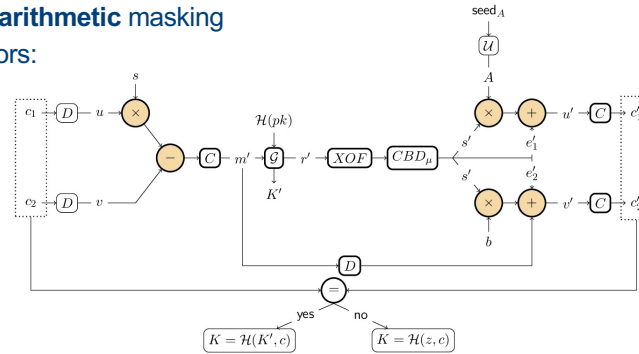
36

36

Polynomial arithmetic

- Easy to protect with **arithmetic** masking
- Small overhead factors:
 - 1.7 to 2.0 (n=2)
 - 2.96 (n=3)

n = sharing factor



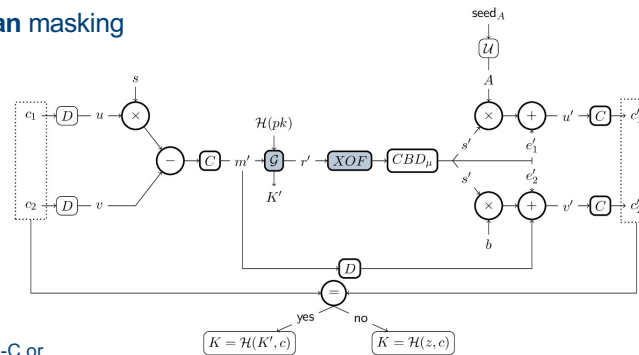
M. VanBeirendonck et al. ACM Journal on Emerging Technologies in Computing Systems 17(2), 25 pages, 2021 [BDK+21]

KU LEUVEN

37

SHA-3

- Protected with **Boolean** masking
- Overhead factors
- 5.9 to 9.26 (n = 2)
- 73.1 (n=3)



Depends if you compare to plain-C or optimized assembly

[Boolean masking: BDPVA10,BBD+16]

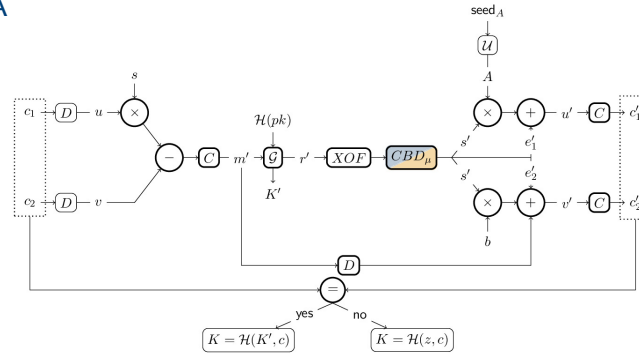
38

KU LEUVEN

38

Centered Binomial sampling

- Mix of A2B and B2A
- Expensive!
- Etc.



39

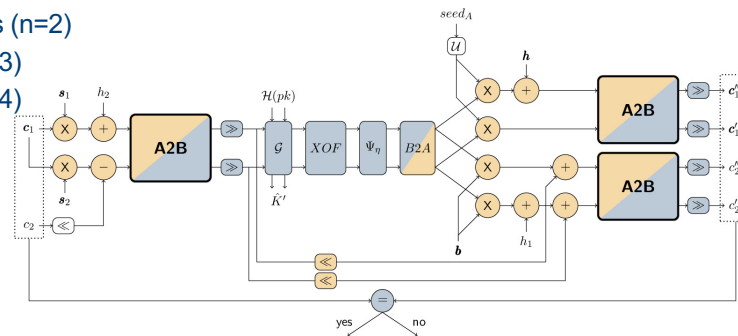
39

One A2B conversion cost (Saber)

Requires bit-slicing

- 55-61 K cycles (n=2)
- 172-206 K (n=3)
- 302-365 K (n=4)

+ randomness



[1] Revisiting Higher-Order Masked Comparison for Lattice-Based Cryptography: Algorithms and Bit-sliced Implementations, D'Anvers J.P., Van Belendonck M., Verbauwhede I., IACR ePrint 2022/110.
 [2] Bitslicing Arithmetic/Boolean Masking Conversions for Fun and Profit with Application to Lattice-Based KEMs, Bronchain O. and Cassiers G., IACR Cryptol. ePrint Arch. 2022: 158 (2022).

40

Masking is expensive

CPU cycles x1000 Scheme	Unmasked	1 st order n=2	2 nd order n=3	3 rd order n=4
Saber	773	3,011 (1x)	5,534 (1x)	8,591 (1x)
Kyber [2]	804	7,716 (2.56x)	11,880 (2.14x)	16,715 (1.94x)
COST	1x	3.9x – 9.6x	7.2x – 14.8x	11.1x – 20.8x
Random bytes		12 KB	42 KB	90 KB

Unmasked Kyber/Saber similar COST

- Masked Kyber more expensive vs Saber
 - Power of two
 - Rounding vs error sampling
- **Masking is expensive AND requires randomness**

Platform: ARM Cortex M4
 Framework: PQM4
 Compiled: arm-none-eabi-gcc
 Version: 9.2.1

41

41

Research challenges for cryptography

- Goal: introduce new research topics, improve existing ones
- Challenge 1: masking is hard in practice
- Challenge 2: masking is expensive
- Challenge 3: Possibilities of PUFs
- Challenge 4: Random number generation
- **Challenge 5: NEW – Fully Homomorphic Encryption**
 - On FPGA
 - On ASIC

42

42

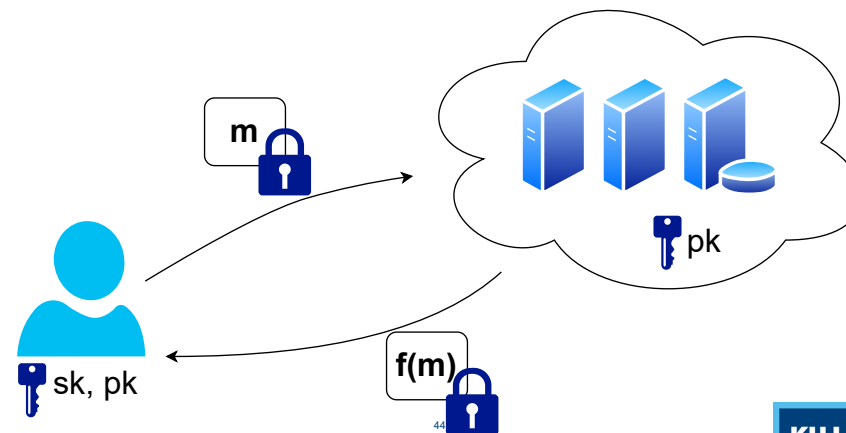
We envision a research agenda in this area that should include, but is not limited to, the following: (1) design for fully secure end-to-end hardware and software solutions that are secure against various forms of attacks on operation, data, and communications; (2) security in the chip design tool chain that would enable end-to-end security solutions to be verified by design; (3) secure hardware supply chain covering chip fabrication, packaging, and system integration; (4) implementation of post-quantum cryptography; (5) implementation of low-power cryptography for secure communications and transactions; and (6) other privacy preserving hardware implementations for processing encrypted data.

43

KU LEUVEN

43

Fully Homomorphic Encryption



44

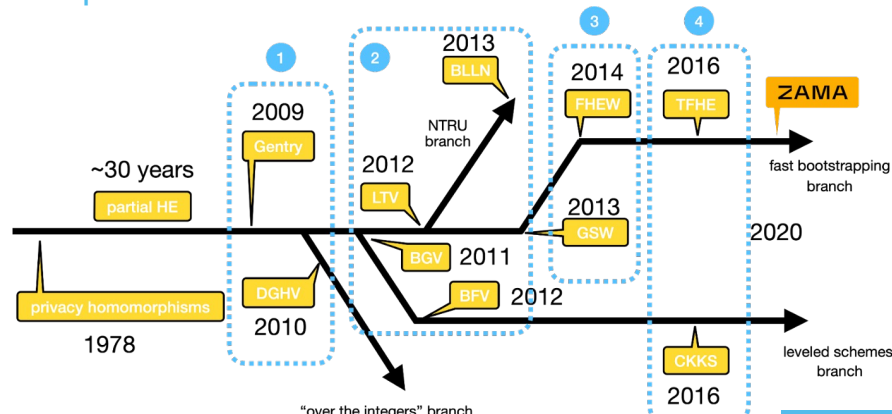
KU LEUVEN

44

Multiple schemes

- Partially homomorphic: Paillier system
- Somewhat homomorphic:
 - Limited number of multiplications
 - Fan-Vercauteren:
- Fully Homomorphic Encryption
 - Unlimited number of multiplications
 - Requires 'bootstrapping'
- Multiple schemes:
 - BFV: Brakerski – Fan – Vercauteren
 - BGV: Brakerski – Gentry – Vaikuntanathan
 - TFHE: Torus Fully Homomorphic Encryption
 - ...

Multiple FHE schemes



[copied from ZAMA website]

Challenge large numbers:

- Experiment 1 [CHES2015] : YASHE (now no longer used, reduced security)
 - Ciphertext size 5MB to 20MB (Polynomial size is 32768 (2^{15}) to 65536 (2^{16}), modulus 1200 to 2500 bits), could evaluate depth of Simon block cipher
- Experiment 2 [TC2018]: HEP CLOUD, FV
 - Ciphertext pair 9.2MB with parameters Polynomial size is 32768 (2^{15}), modulus 1128 bits, depth 36, 85 bits security level.
 - Bottleneck: I/O between FPGA and external memory
- Experiment 3 [TC2020]: HEAWS, FV
 - Cipher text pair 180KB, with parameters Polynomials size is 4096, modulus min 372 (Q), 180 (q), depth 4, more than 80 bits security.
 - Useful for small neural network applications
 - Fits on one FPGA

47

KU LEUVEN

47

DARPA DPRIVE program: in progress

- ▶ Dedicated ASIC acceleration of BGV
 - 150mm² chip in 12nm GF
 - Within 10× of plaintext computation
 - 10,000× faster than software reference
 - Parameter set for 128-bit security
 - Support *bootstrapping*
- ▶ Four teams of researchers
 - Galois, Duality, SRI, and Intel
- ▶ Several phases
 - **Phase 1:** design, implementation and verification of system architecture and IP blocks



Now: phase 2 running, with three teams: Galois, Duality and Intel⁴⁸

KU LEUVEN

48

BGV parameters in DPRIVE

Parameter	Range	Example
Security parameter	N/A	128 bits
Ring dimension N	512 – 65536	65536
Plaintext modulus p^r	≥ 2	127^3
Ciphertext packing ℓ	1 – 65536	64 slots
Max $\log_2(QP)$ for key switching	20 – 1782	1782 bits
Max $\log_2(Q)$ for ciphertext	20 – 1263	1263 bits
Max multiplicative depth L	N/A	31

Ciphertext: 21 MB, Key-switch key: 84 MB

49



49

Hardware acceleration options


50



50

Challenges

- Computational complexity
 - NTT/FFT acceleration
- **Memory**
 - SIZE
 - BANDWIDTH





ASIC (phase 1)

- 150 mm² in 12nm
- Global Foundries
- Memory hierarchy
- 57 – 115 Watt

Cloud FPGA

- Alveo U280 (in 5nm or 7 nm)
- Included into Amazon AWS F1
- Memory hierarchy
- 225 Watt! (cooling)






51

Three experiments – three domain specific processors

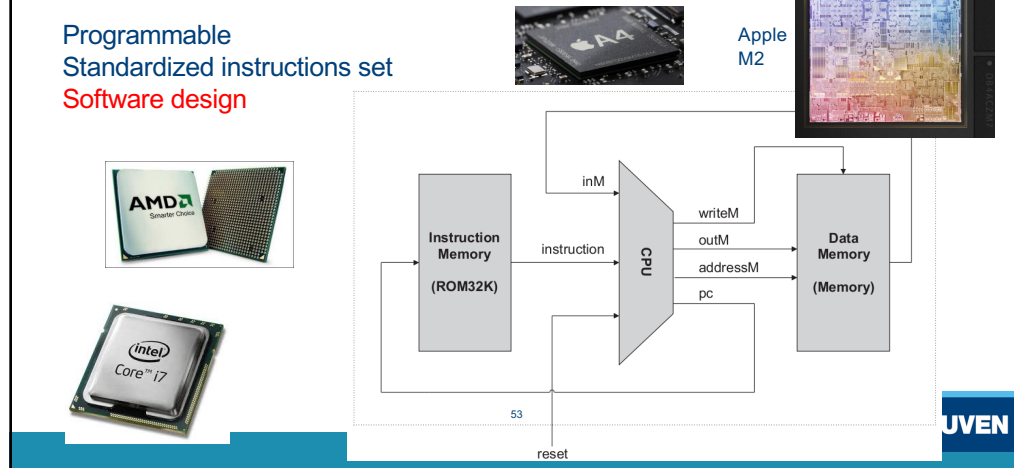
FPGA - HEAWS	ASIC – DPRIVE – BASALISC	FPGA - FPT
<ul style="list-style-type: none"> • BFV – leveled HE • 80 bit security • Shallow depth 	<ul style="list-style-type: none"> • BGV – includes Bootstrap • 128 bit security • DPRIVE constraints 	<ul style="list-style-type: none"> • TFHE • 128/110 bit security • Alveo U280
	<ul style="list-style-type: none"> • NTT acceleration • Residue Number System • Dedicated instruction set • No cache: compile time known 	<ul style="list-style-type: none"> • FFT acceleration • Streaming bootstrap
IEEE TC 2020	IACR 2022/657	IACR 2022/1635



52

Option 1: CPU - General purpose process


Programmable
Standardized instructions set
Software design



53

Option 2: Domain specific processing

- Tightly couple: instruction set extension
 - Register mapped
 - Reuse CPU infrastructure
 - Reuse decode, registers, cache, bus network, etc.
 - Example AES instructions
- Here: FHE specific operations
 - Leads to DOMAIN SPECIFIC PROCESSORS



54

KU LEUVEN

54

Option 3: Domain specific co-processor

- Inside CPU = custom ISA
- Local bus = tight coupled
- Peripheral = loosely coupled

[Picture: P. Schaumont, "A practical introduction to Hardware/Software Codesign", 2nd ed

KU LEUVEN

55

Hardware technology: FPGA versus ASIC

FPGA

- Field Programmable Gate Array

Xilinx XC4000ex (OLD!)

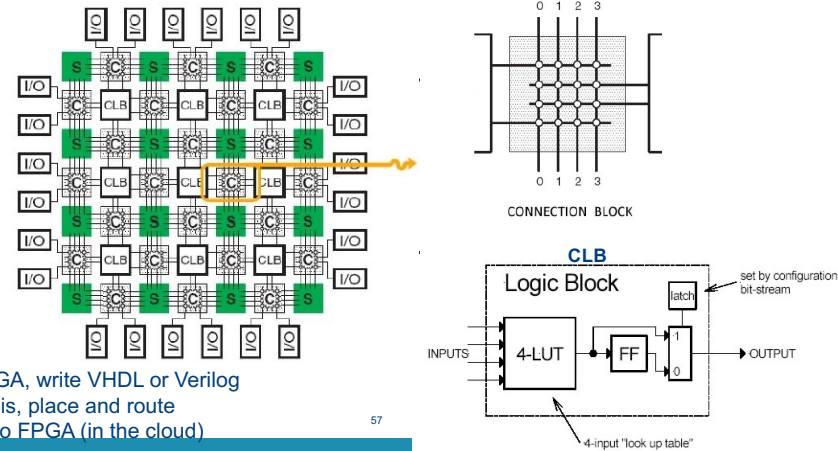
ASIC

- Application Specific Integrated Circuit

KU LEUVEN

56

FPGA: Program look-up tables and interconnect



57

First experiment: FPGA Acceleration of BFV on Amazon cloud

58

FPGA Memory Resources (Alveo U280)

BRAM – 9 MB

LOCAL ON CHIP

URAM – 33 MB


LOCAL ON CHIP

HBM – 8 GB

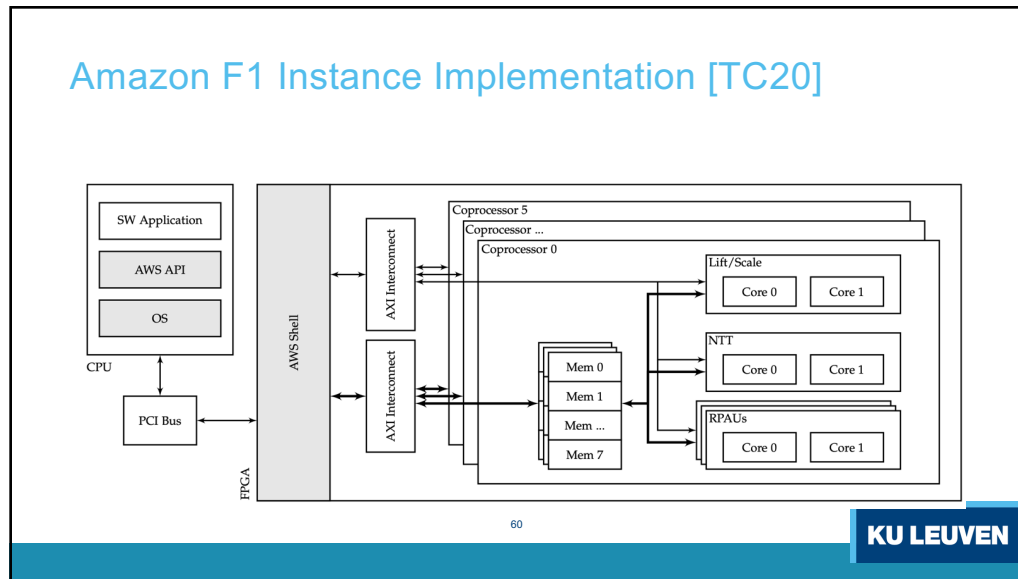
IN PACKAGE, 3D

DDR – 32 GB

ON BOARD



59



60

Performance of Homomorphic Multiplication

- Each multiplication takes 4.34 ms.
- The overhead of a ciphertext transfer is 0.11 ms.
- A single coprocessor achieves 230 multiplications per second.
- **Six** coprocessors running in parallel achieves 613 multiplications.

61

KU LEUVEN

61

Comparison

- Achieve 613 homomorphic multiplications per second
- Compared to CPU
 - 13x speedup w.r.t. a highly optimized software on Intel i5 processor, 1.8 GHz
- To GPU on Amazon cloud, **5 times more work for half price and lower power!**

Compute: Amazon EC2 Instances:

Description	Instances	Usage	Type	Billing Option	Monthly Cost
1 FPGA -> 2000 Mult	1	100 % Utilized/Mc	Linux on f1.2xlarge	On-Demand (No Co	\$ 1207.80
1 GPU -> 388 Mult	1	100 % Utilized/Mc	Linux on p3.2xlarge	On-Demand (No Co	\$ 2239.92

14.02.2019 - AWS Simple Monthly Calculator: <https://calculator.s3.amazonaws.com/index.html>
17.06.2020

62

KU LEUVEN

62




Second experiment: ASIC Acceleration of BGV

Darpa DPRIVE Basalisc project

KU LEUVEN

63



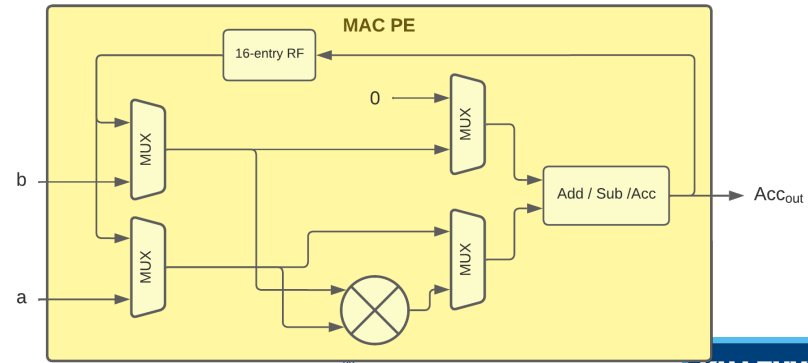
BASALISC Memory Hierarchy

MAC Acc – 8 KB	
MAC Register File – 128 KB	
Cipher Text Buffer – 64 MB	ON CHIP, Cipher Text Buffer CTB fits 3 ciphertext pairs
DDR – 256 GB	One Key switch 84MB does not fit ON BOARD

KU LEUVEN

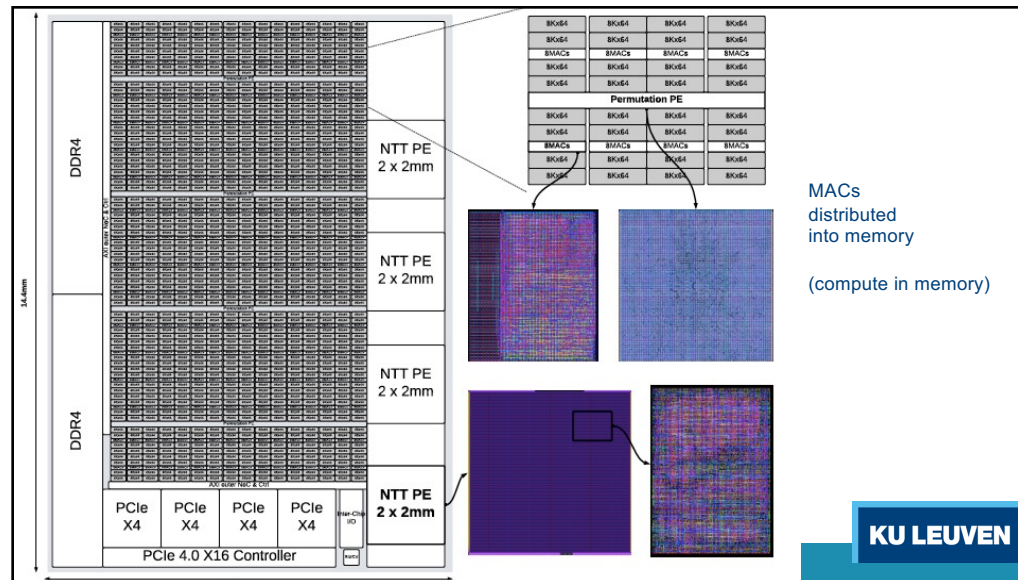
64

2048 x 32 bits Modular Multiply – Accumulate unit PE



KU LEUVEN

65




MACs distributed into memory
(compute in memory)

KU LEUVEN

66

NTT units with conflict free access to Cipher Text Buffer

- 1 Radix-256 butterfly
 - 65536-point NTT with 2 passes
 - 2 Twiddle-factor factory unit
 - 3 Conflict-free data layout
- 32 Tb/s NTT throughput**



Example radix-4 butterfly

67

67

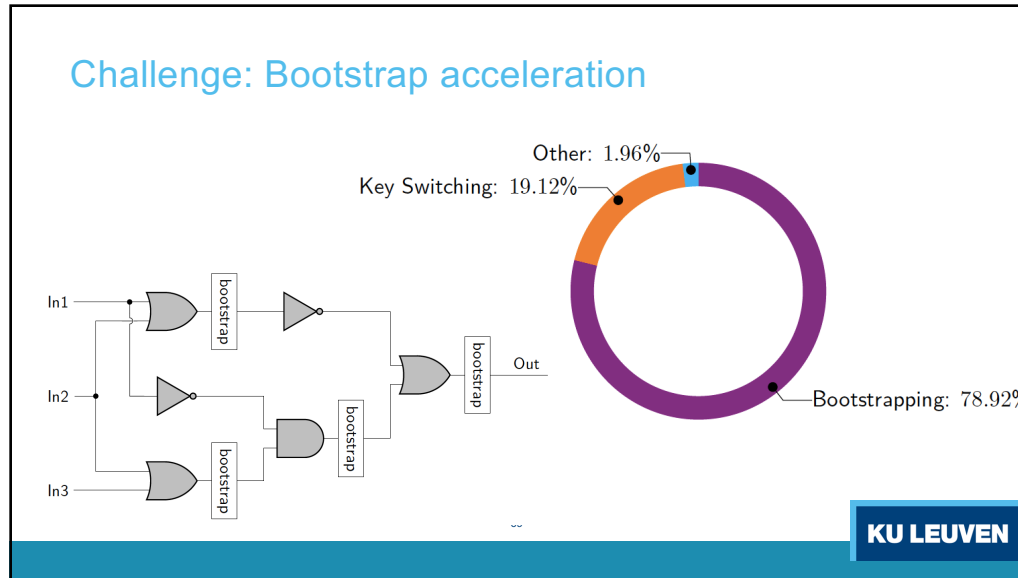


Third experiment: FPT FPGA Fixed Point Accelerator for TFHE Torus Fully Homomorphic Encryption

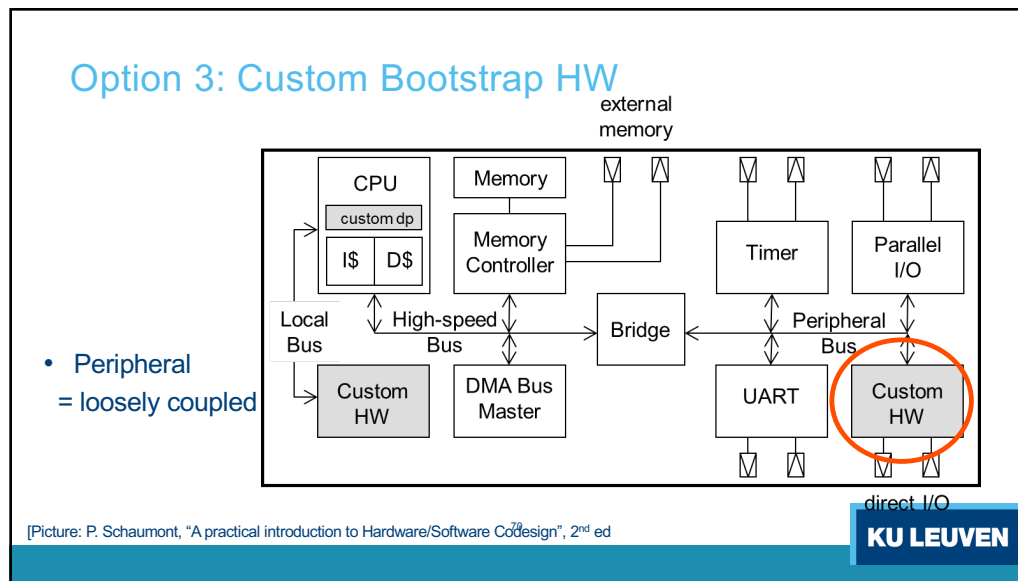
ERC Advanced Grant Belfort, FWO

68

68

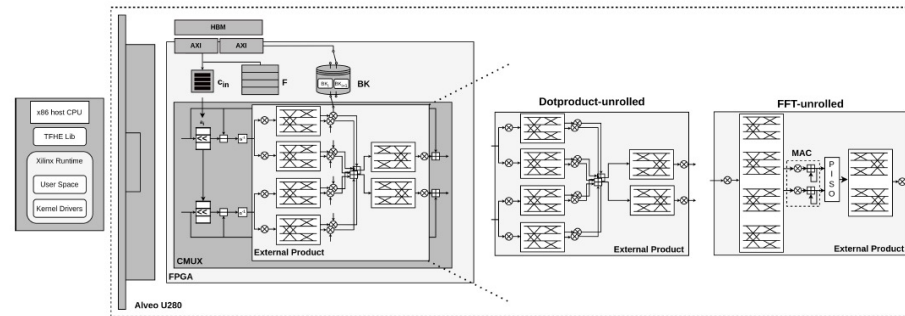


69



70

FPGA: Bootstrap FFT accelerator



71

71

Results

	LUT / FFs / DSP / BRAM	<i>Clock frequency</i> f (MHz)	<i>Latency</i> l (ms)	<i>Throughput</i> TP (PBS/ms)
• FPGA	FPT 595K / 1024K / 5980 / 14.5Mb	200	0.58	25.0
	YKP 842K / 662K / 7202 / 338Mb 442K / 342K / 6910 / 409Mb	180 180	3.76 1.88	3.5 2.7
• ASIC	MATCHA 36.96mm ² 16nm PTM	2000	0.2	10
• CPU	CONCRETE Intel Xeon Silver 4208	2100	32	0.03
• GPU	cuFHE NVIDIA GeForce RTX 3090	1700	9.34	9.6

72

72

Conclusions – lessons learned

- “Provable secure masking” does not mean secure: theory and practice are different.
 - Practical evaluation in the lab of theoretical security is a must
 - Papers should include artifact evaluation.
- Masking and especially higher order masking are expensive, orders of magnitude
 - Less stringent, more realistic models
 - Reduce randomness requirements
- Fully Homomorphic Encryption
 - New research topic for HW acceleration

73

KU LEUVEN

73

(e) Semiconductors and System Security

Criminal and state-sponsored cyber-attacks pose increasing threats to the United States. To enable the implementation of secure systems, every aspect of the system must be considered including sensors, data converters, computing, memory, storage, and communications, while providing robustness against side-channel attacks and ensuring security of supply chains. There is a tremendous opportunity for the design of secure semiconductor chips. To maximize effectiveness, security must be pursued as an integral part of design, not as an add-on after the chip is designed.

Academia, industry and government stakeholders have an opportunity to standardize a trusted approach for systems implementation. The specific opportunity is to bring together algorithm and software/systems designers with chip designers in a center of excellence to develop the next generation of secure systems. Although open-source security approaches are the best for innovation and transparency, they remain unpalatable for the industry. We must address this reluctance in a way that enables the United States to continue to be the global leader in standardized security approaches.

We envision a research agenda in this area that should include, but is not limited to, the following: (1) design for fully secure end-to-end hardware and software solutions that are secure against various forms of attacks on operation, data, and communications; (2) security in the chip design tool chain that would enable end-to-end security solutions to be verified by design; (3) secure hardware supply chain covering chip fabrication, packaging, and system integration; (4) implementation of post-quantum cryptography; (5) implementation of low-power cryptography for secure communications and transactions; and (6) other privacy preserving hardware implementations for processing encrypted data.

VEN

74