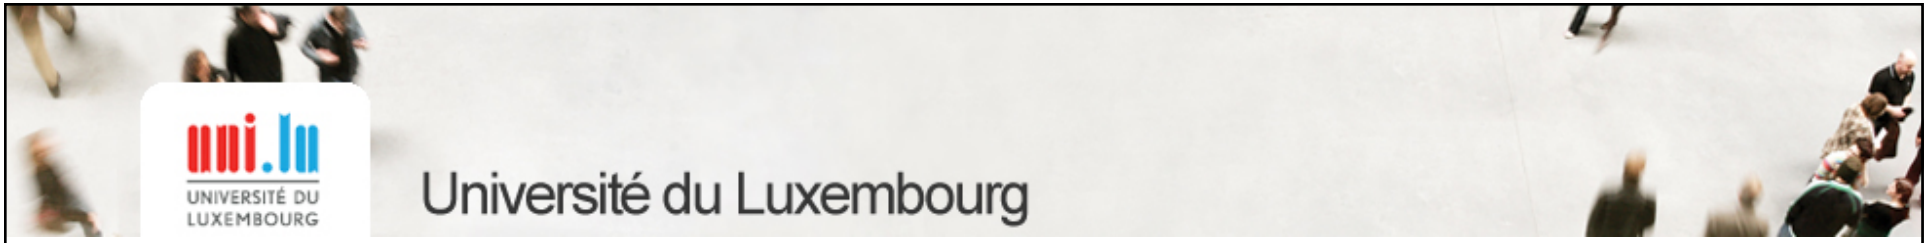

CAESAR: Cryptanalysis of the Full AES Using GPU-Like Hardware

Alex Biryukov and Johann Großschädl

Laboratory of Algorithmics, Cryptology and Security
University of Luxembourg

SHARCS 2012, March 17, 2012

Motivation

- Cryptanalytic attack on the full AES
  - Suppose complexity of < 2^{100} computations
  - AES-256: related-key cryptanalysis ($2^{99.5}$)
  - AES-128: TMK trade-off (e.g. 2^{96})
- Using special-purpose hardware
  - Is it feasible?
  - Totally infeasible?
  - How would you design such hardware?
  - Where are the bottlenecks?

Outline

- Cryptanalytic attacks on AES
- GPU-like AES processor for cryptanalysis
- Memory and storage
- CAESAR supercomputer
- Time and energy
- Outlook into the future

History

- Special-purpose hardware for cryptanalysis
  - Cryptanalysis of Enigma and Lorenz
  - Quasimodo for factoring
  - DES-cracker (EFF)
  - TWINKLE and TWIRL
  - COPACOBANA
- Software
  - Cell Processor
  - GPUs

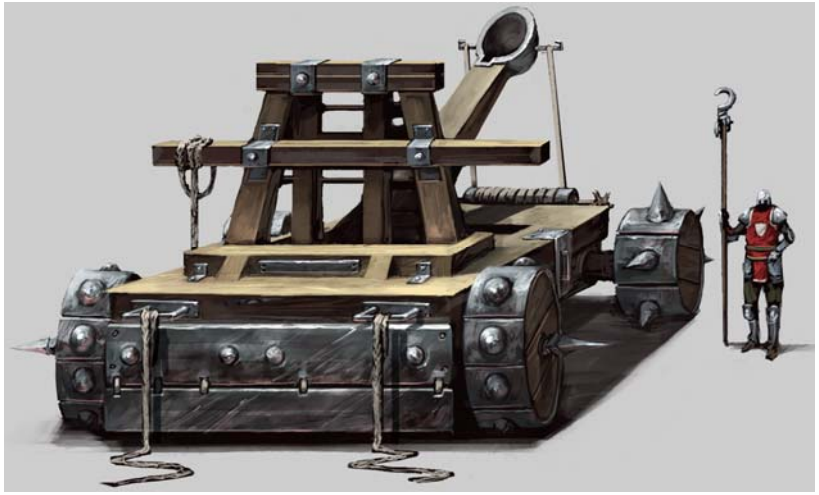

Cryptanalytic Attacks on AES

Cipher	Attack/Result	Rounds	Data	Workload	Memory	Reference
AES-128	Multiset	6	2^{33}	2^{70}	2^{32}	Daemen (2002)
	Collisions	7	2^{32}	2^{128}	2^{80}	Gilbert (2000)
	Partial sum	6	2^{35}	2^{44}	2^{32}	Ferguson (2000)
	Partial sum	7	$2^{128} - 2^{119}$	2^{120}	2^{32}	Ferguson (2000)
	Boomerang	6	2^{71}	2^{71}	2^{33}	Biryukov (2004)
	Impossible diff.	7	$2^{112.2}$	$2^{117.2}$	2^{109}	Lu (2008)
	Boomerang - RK	7	2^{97}	2^{97}	2^{34}	Biryukov (2009)
AES-192	Rectangle - RK	9	2^{64}	2^{143}	?	Gorski (2008)
	Rectangle - RK	10	2^{125}	2^{182}	?	Kim (2007)
	Boomerang - RK	12	2^{116}	2^{169}	2^{145}	Biryukov (2009)
AES-256	Rectangle - RK	10	2^{114}	2^{173}	?	Biham (2005)
	Subkey Diff.	10	2^{48}	2^{49}	2^{33}	
	Differential - RK	14	2^{131}	2^{131}	2^{65}	Biryukov (2009)
	Boomerang - RK	14	$2^{99.5}$	$2^{99.5}$	2^{78}	Biryukov (2009)

Related-Key Boomerang Attack

- Example: AES-256
  - Attacker knows relation betw. 4 secret keys
  - Time and data complexity: $2^{99.5}$
  - Memory complexity: 2^{78}
- Motivation for using RK attack
  - No practical impact due to reliance on related keys
  - However, future single-key attacks may have similar structure and requirements

Time-Memory-Key Trade-Off

- Example: AES-128
  - Fixed plaintext encrypted under 2^{32} keys
  - 2^{96} off-line pre-computation
  - 2^{80} on-line computation with 2^{56} memory

$$N^2 = T(MD_k)^2$$

Attack	Data Type	Keys (Data)	Time	Memory	Preprocessing
BS TMD	FKP	2^{8}	2^{120}	2^{60}	2^{120}
BS TMD	FKP	2^{20}	2^{100}	2^{58}	2^{108}
BS TMD	FKP	2^{32}	2^{80}	2^{56}	2^{96}
BS TMD	FKP	2^{43}	2^{84}	2^{43}	2^{85}

GPU-like AES Processor

- GPU architecture
  - Homogenous multi-core system local cache
  - Much higher performance than a CPU

Normal CPU:
4-8 ALUs

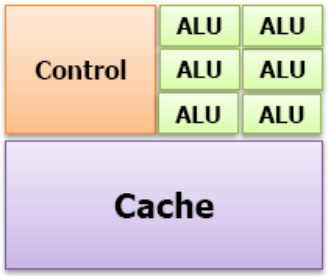

GPU: several
100 ALUs

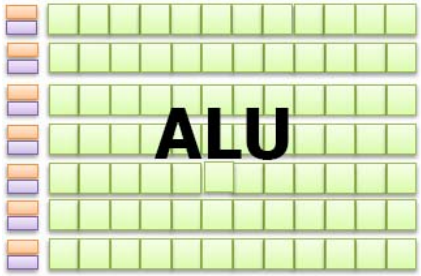

- Optimized for cryptanalysis of AES
  - Replace CUDA cores by AES cores
  - Data/keys either constant or generated on-chip

AES Core

- Optimized for high throughput
  - Loop unrolling: 128-bit datapath for each round
  - Do not need to support a mode of operation
  - Inner-loop and outer-loop pipelining
- Example: Hodjat-Verbauwhede (2003)
  - Architecture as above
  - New plaintext each clock cycle
  - Every round takes 4 clock cycles
  - Latency of 41 cycles

Hodjat's AES Processor

- Performance
  - 0.18 µm standard-cell library (2003)
  - Can be clocked with 606 MHz (pipelining)
  - Throughput: 77.6 Gbit/s
- Silicon area
  - 473k gates (for 10 rounds)
  - 660k gates (for 14 rounds for 256 bit keys)

NVIDIA GT200b

- Main Characteristics
  - 240 “shader” cores, each up to 3 FLOP/cycle
  - 1476 MHz, 350M gates (470 mm^2)
  - 1000 GFLOPS (10x faster than Intel Core i7)

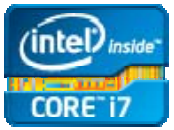

- GT200b is well-known in crypto community
  - But today not state-of-the-art anymore

GPU-like AES Processor

- 500 AES cores based on Hodjat's design
- Each 660k gates, i.e. 330M gates in total
- Clocked at 2.0 GHz on 55 nm TSMC technology (requires better cooling)
- Throughput: $500 \times 2 \cdot 10^9 = 10^{12}$ AES ops/sec

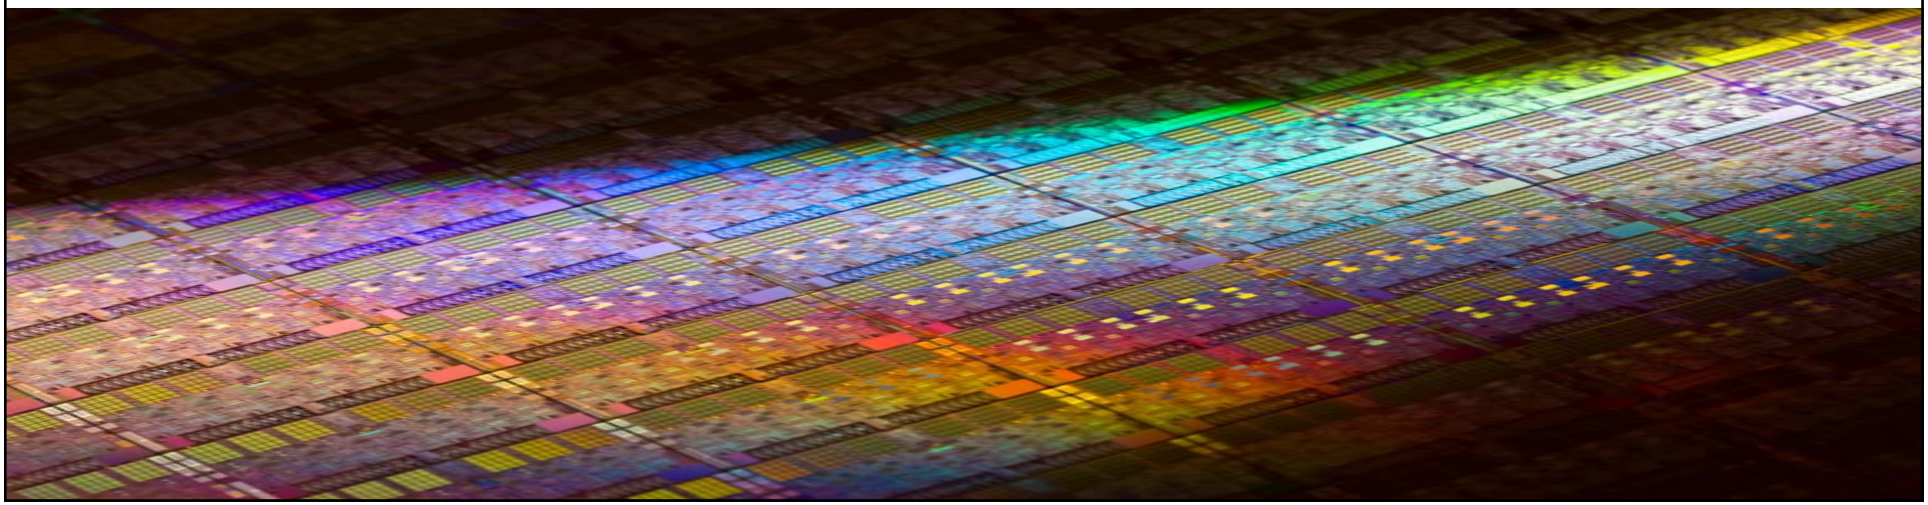

I/O Requirements

- Data (i.e. plaintexts) and keys
  - Bandwidth no problem for data and key
  - RKC: no need for key-agility (4 fixed keys); plaintext generated on chip (using counter)
  - TMK: plaintext is constant; keys can be generated on chip
- Ciphertexts
  - Only “few” ciphertexts are actually stored
  - Transfer rate 2^{21} slower than processing rate

Storage Requirements

- RKC: 2^{78} bytes
  - $3x10^{10}$ harddisks of 10 TBytes each
  - Bottleneck today, but feasible in 5-10 years
  - Current state-of-the-art: 4 TB, 250$ retail price
  - 100 TB @ 100$ in 5-10 years: 300 bln $
- TMK: 2^{61} bytes
  - 92 mln $ now
  - 2.3 mln $ in 5-10 years

CAESAR Supercomputer

- Assumptions about adversary
  - 1 trillion (10^{12}) US$ for chip production
  - Roughly US national defense budget in 2010
  - US budget deficit >1.4 trillion US$ in 2009
- **C**ryptanalytic **AES AR**chitecture
  - Hypothetical supercomputer to break AES
  - 3×10^{10} GPU-like AES processors (30$/proc.)
  - 3×10^{22} AES operations/sec
  - 10 TB storage attached to each AES processor

Computation Time

- 3×10^{22} AES operations/sec
  - We assume attack time being determined by AES computations and not access to storage
- RKC attack
  - 1 year for $2^{99.5}$ AES ops (lower bound)
- TMK attack
  - Pre-computation (2^{96} AES ops): 1 month
  - Online phase (2^{80} AES ops): negligible
  - 1/10th of budget: 1 year pre-comp, 8 min online

Production of Chips

- High capacity fab: 300,000 wafers/month
- About 100 AES processors per wafer
- We need 83 fabs (1 bln US$ each); time to build a fab: 18 months; fabs work 1 year

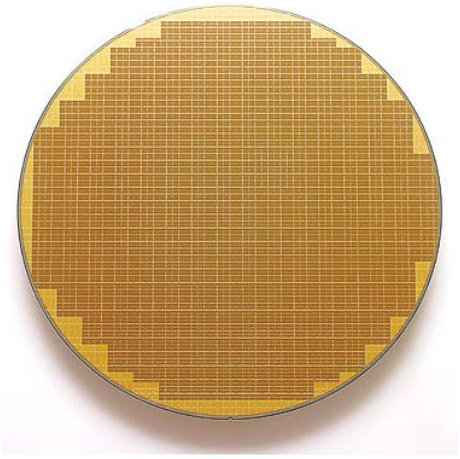

Energy

- 135 W per AES processor
- For $3x10^{10}$ processors – 4 TW
- US power consumption per year: 3.34 TW in 2005
- Water cooling
- Energy seems to be the main *bottleneck*

Energy: Impact of Moore's Law

- Shrinking transistor sizes:
  - From 55 nm in 2009 to 7 nm by 2020
- Historical example:
  - First 1 TFLOPS supercomputer was ASCI Red by Intel (1997) for Sandia Labs
  - 10,000 Pentium Pro (333MHz): **500kW**
  - Now a single GT200b 1 TFLOPS @ **100 W**
  - Factor 5000 in 13 years

Summary

- **C**ryptanalytic **AES AR**chitecture

One AES engine:	660K gates	2GHz clock speed	
One AES processor:	500 AES engines	10^{12} AES ops/s	30 US$
CAESAR supercomputer:	$3 \cdot 10^{10}$ AES processors	$3 \cdot 10^{22}$ AES ops/s	1 trillion US$
AES chip production:	83 high capacity fabs	approx. 1 year	83 bln US$
Power consumption:	135 W per processor	4 TW = $4 \cdot 10^{12}$ W	
RKC Attack:	$9 \cdot 10^{29}$ AES ops	approx. 1 year	
Storage RKC:	2^{78} bytes		300 bln US$
TMK Attack (2^{32} targets):	$0.8 \cdot 10^{29}$ pre-computation	30.6 days	
TMK Attack (2^{32} targets):	10^{24} ops per AES key	negligible	
Storage TMK:	2^{61} bytes		92 mln US$

Outlook into Future

- Cryptanalytic breakthroughs?
- Moore's law (10 more years)
- Computers based on spin (spintronics) or optical computers
- New storage technologies
  - Thermally assisted recording (10TB/inch)
  - Further miniaturization (12 atoms/strorage cell)
  - Quantum holography
  - 3D optical storage (1TB DVD with 100 layers)

Conclusions

- Hypothetical supercomputer CAESAR
  - TMK on AES-128 with 2^{32} targets is well within reach of current VLSI technology
  - For $2^{99.5}$ time/data attack with 2^{78} memory: memory complexity and power consumption are main bottlenecks, but not execution time
- Recommendations
  - Focus on attacks with time complexity of up to 2^{100}, but memory complexity of less than 2^{70} (and as little data as possible)

Questions?