Radboud University

A5/3 make or break: A massively parallel FPGA architecture for exhaustive key search

Konstantina Miteloudi, Lejla Batina, Nele Mentens

CHES, September 2025

Security in 2G (GSM) mobile networks

- A5 family of algorithms:
  - A5/1 and A5/2 are already broken
  - A5/3 is a stream cipher.
    - It is based on KASUMI block cipher.
    - KASUMI uses 128-bit key.
    - A5/3 expands a 64-bit session key K_c to 128-bit KASUMI key by concatenation, $K_c||K_c$.
    - This 2^{64} effective key space makes A5/3 a good candidate for brute-force attacks.
  - A5/4 same with A5/3 but with 128-bit session key.

2G networks still in use

- 2G and 3G are phasing out, but:
  - with long timetables (e.g. UK has set a deadline of 2033), and
  - it is expected that 2G will remain a legacy network for a long time.[a]
  - e.g. Malaysia shutdown 3G in 2021, but kept 2G alongside the 4G and 5G networks.
- 2G still in use for:
  - coverage (rural areas), Machine to Machine communication (M2M) (Sensors), Emergency calls services (eCall) and more.[b]

[a] GSMA. The state of mobile internet connectivity, 2023 and GSMA. NG.121 - 2G-3G sunset guidelines version 2.0, 12 2024

[b] Body of European Regulators for Electronic Communications (BEREC) : Report on practices and challenges of the phasing out of 2G and 3G, 2023

2G coverage in Kuala Lumpur

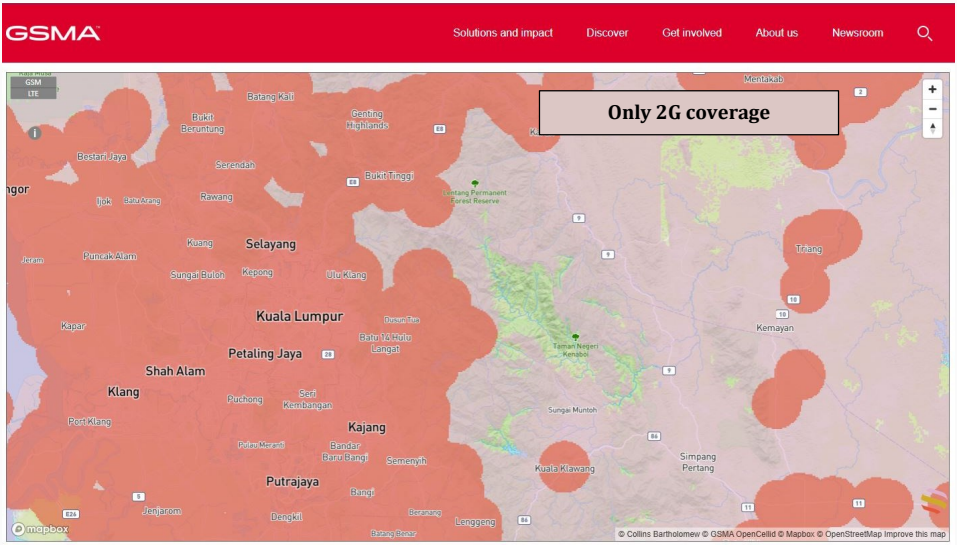

Feasibility of Exhaustive Key Search on A5/3

- *Is A5/3 practically breakable by exhaustive key search using contemporary hardware?*
  - How long will it take?
  - How much will it cost?

Computing power : then vs now!

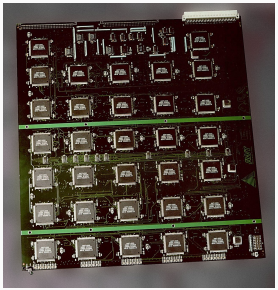

1998: Deep Crack
(ASICs)

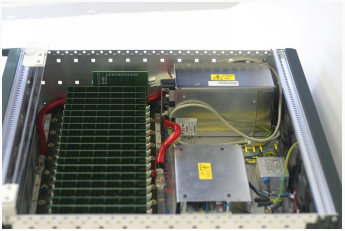

2006: COPACOBANA
(120 Spartan-3 FPGAs)

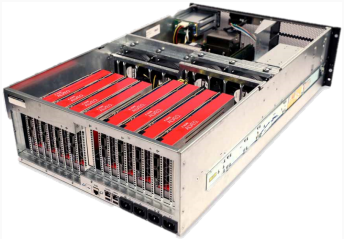

Cloud-FPGA Server
(8 Alveo U250)

GSM A5/3 encryption/decryption

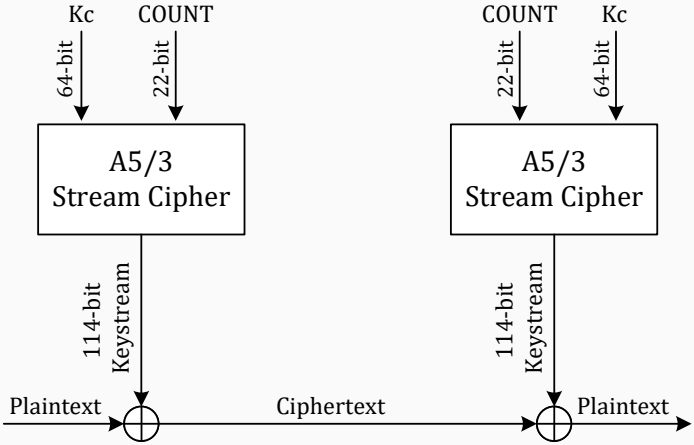

Input *COUNT* obtained from the Frame Number of the Time-Division Multiple Access (TDMA).

Attack Scenario

- Attack scenario : known plaintext-ciphertext and parameters of IV
  - Feasible to obtain these data from the GSM network (Avoine et al. – CRYPTO 2024)
  - therefore known *keystream* $= p \oplus c$ and the XOR step is omitted.
- Problem is reduced to:
  - how fast can we generate the keystreams for all possible keys?
- **NOTE:** Attacks on KASUMI are not applicable on A5/3.

A5/3 specifications

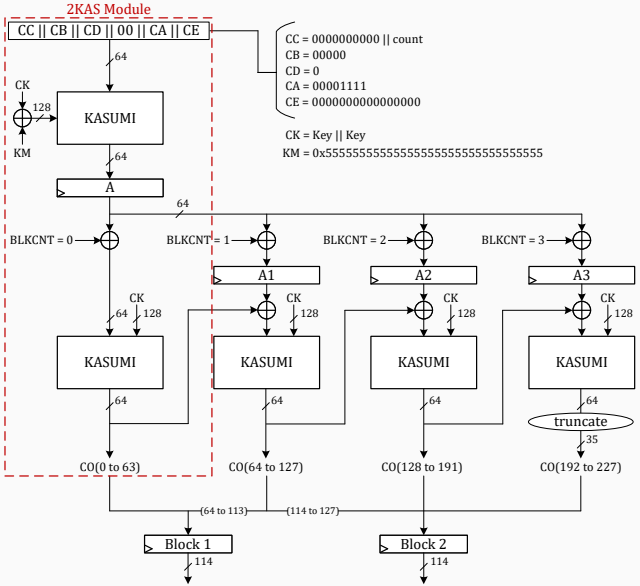

KASUMI block cipher

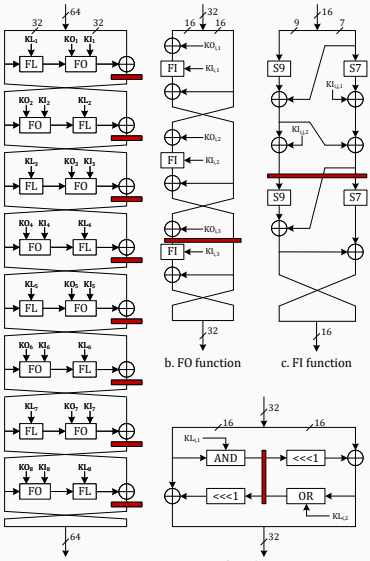

Hardware Architecture (Core)

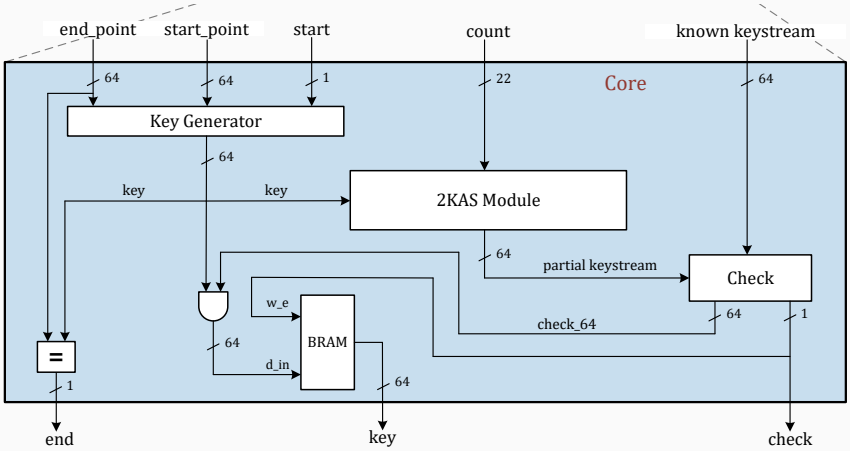

c. Basic computation core

Hardware Architecture (FPGA Alveo U250 and Kernels)

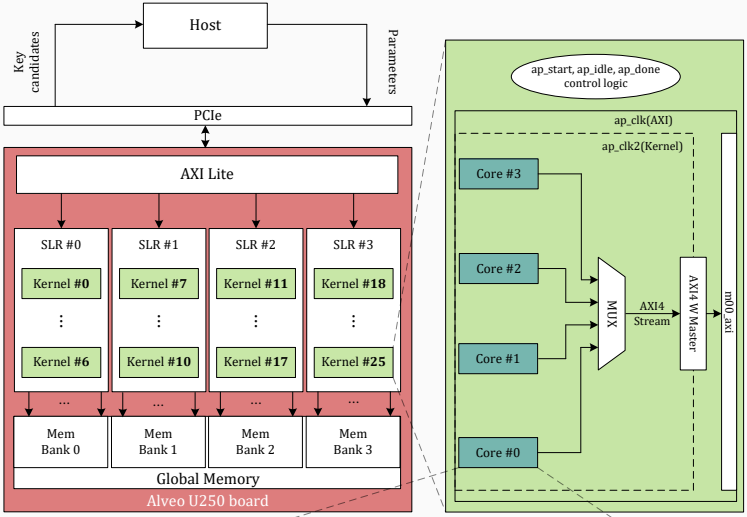

a. Overall Architecture

b. RTL Kernel block

Implementation Results

- Utilization

Alveo U250		KASUMI	1-Core	104-Core
Resources	**Available**	**Used (%)**	**Used (%)**	**Used (%)**
LUT	1 726 216	4 781 (0.28)	10 074 (0.58)	1 193 426 (69.14)
LUTRAM	790 200	960 (0.12)	2 206 (0.28)	192 647 (24.38)
FF	3 456 000	2 928 (0.08)	6 415 (0.19)	969 246 (28.05)
BRAM	2 688	N/A	1 (0.04)	606 (22.54)
DSP	12 288	N/A	N/A	13 (0.11)

- Timing
  - 104 cores at clock frequency of 496.7 *MHz* (2.013 *ns*)
  - $2^{35.59}$ keys/second (51.72 billion keys/second) per Alveo U250 board

Scalable system overview (Level 1 and 0)

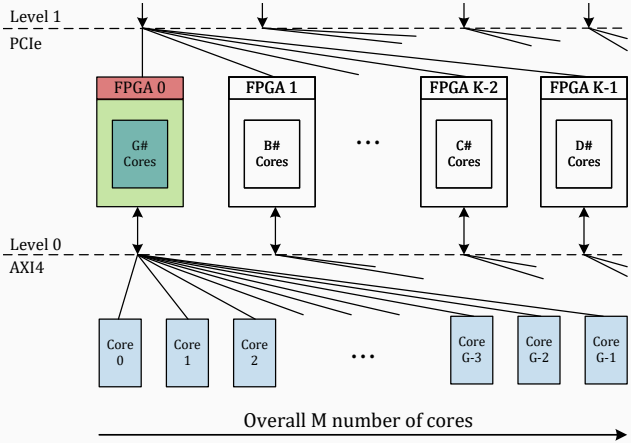

Scalable system overview (Level 2)

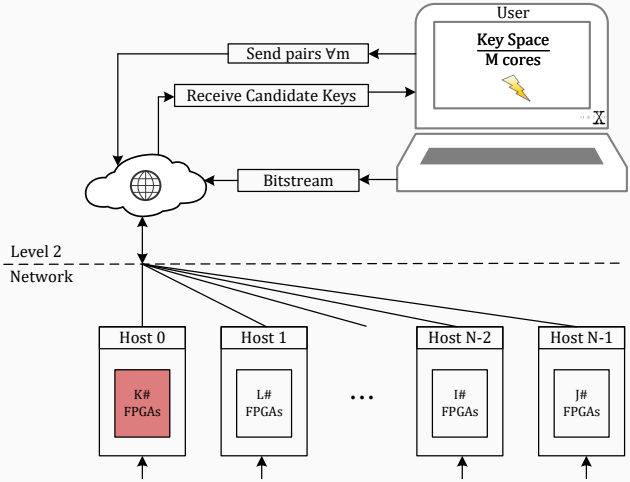

Execution & Expected Time

- The execution time of exhaustive key search is:

$$\textbf{Execution Time} \approx (\mathbf{K} * \mathbf{t})/(\mathbf{c} * \mathbf{b})$$

  - with $K = 2^{64}$ the number of possible keys,
  - $t = 2.013$ *ns* the time for 1 key evaluation,
  - $c = 104$ the number of cores and b the number of boards.

- The expected time of the attack:

$$\textbf{Expected Time} \approx \textbf{Execution Time}/\mathbf{2}$$

  - On average we expect to recover the key in half way of the search.

Time and Cost Analysis

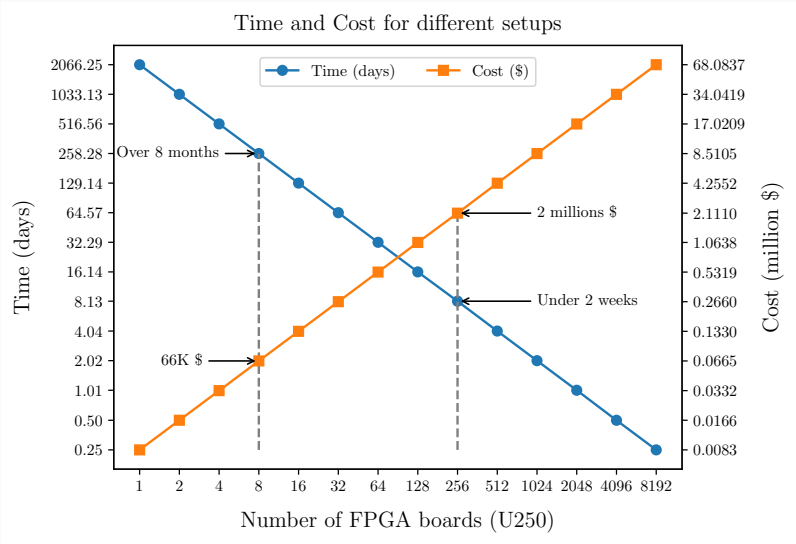

- Unit price $8311

Conclusions

- Whole key space 2^{64} in 1 day with cost of hardware \$34 million.
  - A5/3 should not be considered secure anymore.
- Mounting the attack from the cloud:
  - Microsoft Azure instances (NP series) : the cost per attack is \$81 823.
- All legacy systems using 2G mobile communication must be evaluated:
  - particularly in critical industrial infrastructures (sensors and actuators).
- This work can be used as a framework for accelerating more complicated cryptanalytic attacks.

Protect your own phone

- Many phones still allow **2G fallback by default**.
- Some phones allow users to **disable 2G service manually** in network settings.
- Simple but effective **mitigation step** for end users.

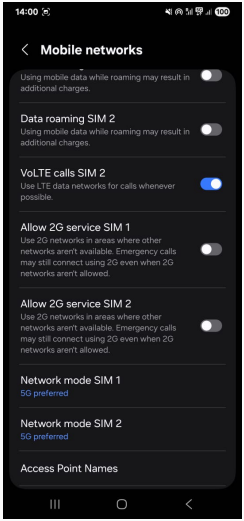

Terima Kasih :)

FPGA vs ASIC

FPGA (U250)

- Implementation
- 104 cores @ 496.7 MHz
- Throughput: $2^{35.59} \approx 51.7$B keys/s
- Cost: \$8,311
- Efficiency: 1.61×10^{-7} \$/key/s

ASIC (16mm^2, GF 22FDX)

- Synthesis
- 240 cores @ 1.1 GHz
- Throughput: $2^{37.94} \approx 263.7$B keys/s
- Cost: \$252,160 (excl. NRE)
- Efficiency: 9.56×10^{-7} \$/key/s

ASIC is $\sim$5$\times$ faster, but FPGA is $\sim$6$\times$ more cost-efficient.

ASIC cost is for a Multi-Project Wafer (MPW) run, fabrication only.
Excludes NRE, packaging, and testing — real project cost would be much higher.

FPGA vs GPU

FPGA (U250, our work)

- 104 cores @ 496.7 MHz
- Throughput: $2^{35.59} \approx 51.7$B keys/s
- Unit price: \$8,311

GPU (RTX 4090, Tez24, KLEIN-64)

- Throughput: $2^{35.40} \approx 45$B keys/s
- Unit price: \$1,929

GPU (RTX 3090, ACC+24, TMTO)

- Throughput: $2^{31.47} \approx 2.9$B keys/s
- Unit price: \$1,000

- FPGA results are **for A5/3 directly** (2 KASUMI blocks).
- RTX 4090 result is for **KLEIN-64**, a $\sim 10\times$ lighter cipher, straight comparison is misleading.

- RTX 3090 result is **17× slower** than FPGA.
- TMTO approach: faster exploitation, but massive precomputation.
- For similar (total) cost ($\approx$\$4.4M), our FPGA approach can search the full 2^{64} keyspace in **8 days**.

[Tez24] targets a lighter algorithm, and [ACC+24] focuses on KASUMI with costly TMTO. Our FPGA work evaluates A5/3 directly, and can reduce precomputation time if combined with [ACC+24].