



HAL
open science

Emulating Side Channel Attacks on gem5: lessons learned

Lilian Bossuet, Vincent Grosso, Carlos Andres Lara-Nino

► **To cite this version:**

Lilian Bossuet, Vincent Grosso, Carlos Andres Lara-Nino. Emulating Side Channel Attacks on gem5: lessons learned. 8th IEEE European Symposium on Security and Privacy (EuroSP2023), Jul 2023, Delft, Netherlands. 10.1109/EuroSPW59978.2023.00036 . hal-04165220

HAL Id: hal-04165220

<https://hal.science/hal-04165220>

Submitted on 18 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Emulating Side Channel Attacks on *gem5*: lessons learned

Lilian BOSSUET, Vincent GROSSO, and Carlos Andres LARA-NINO
Université Jean Monnet Saint-Étienne, CNRS, Institut d'Optique Graduate School,
Laboratoire Hubert Curien UMR 5516, F-42023,
SAINT-ETIENNE, France.
carlos.lara@univ-st-etienne.fr

Abstract

Side channel attacks (SCA) have the potential of disrupting the trust of the users on computing platforms and cryptographic algorithms. The main challenge in the design of countermeasures against such threats is that an evaluation of their effectiveness can only be performed after they have been implemented. By that point, significant resources would have been invested in the creation of a prototype. Moreover, the large volume of combinations from all the potential target algorithms and computing systems complicates a systematical analysis. It is necessary to find strategies to simplify and systematize the study of SCAs and their countermeasures. *gem5* is a cycle-accurate simulator which offers the possibility to emulate a broad range of computing architectures. Beyond the functional verification, this tool computes multiple physical statistics from the simulated system. In this paper, we discuss the lessons learned from using *gem5* to simulate SCAs on an ARM system. Our work shows that while there is a correlation between the data and the reported statistics, there are significant challenges that must be addressed to improve the use of *gem5* for the emulation of physical phenomena.

Keywords: Power Covert Channels, Power Analysis, Gem5.

Please cite as:

```
@InProceedings{BGL23,  
  title = {{Emulating Side Channel Attacks on gem5: lessons learned}},  
  author = {Bossuet, Lilian and Grosso, Vincent and Lara-Nino, Carlos Andres},  
  booktitle = {Proceedings of the 2023 IEEE European Symposium on Security and Privacy  
  Workshops (EuroS&PW)},  
  pages = {287--295},  
  year = {2023},  
  publisher = {IEEE},  
  location = {Delft, Netherlands},  
  doi = {10.1109/EuroSPW59978.2023.00036},  
  isbn = {979-8-3503-2720-5}}
```

1 Introduction

Modern computing systems (like the IoT, wearable devices, and autonomous vehicles) exhibit characteristics that distinguish them from conventional computers. They run in large, unstructured environments to bring services closer to the user. They are networked and collaborate with each other to solve complex tasks with a limited investment of resources per device. They are ubiquitous and heterogeneous. But most importantly, they manage large volumes of data which may require to be protected from unauthorized access or modification. Cryptographic algorithms are the tools of choice for supplying security services such as confidentiality and integrity to the data. They may take different forms and offer distinct performance and cost tradeoffs as to be of use for many computing systems. But all of them, from a theoretical point of view, base their security on hard computing problems. Finding the correct solution [Ber05] would be harder than looking for a particular grain of sand on earth [Kru12].

Under formal security assumptions [BR05] it is necessary to consider that a potential attacker may gain access to the forward transformation of a cryptosystem (chosen plaintext attack) and even its inverse function (chosen ciphertext attack). To regard the system as secure, it must resist any potential attacks when these oracles are available to the adversary. That is, the attacker shall not gain any other information from the queries performed to the functional oracles. However, in practice, there are more sources of knowledge which may give the adversary an advantage to compromise the security of the system. Any computing device is, by nature, an agglomeration of physical phenomena. So goes the phrase “we have tricked a rock into thinking.” Consequently, there are magnitudes which can be seen and quantified when the system runs. Some of these measurements are bound to be correlated with the data being processed by the platform [KJJ99]. If the system happens to perform cryptographic operations, these will also affect the different magnitudes of the device: supply power, electromagnetic emanation, clock frequency, heat dissipation, etc. Thus, it is said that data is *leaked* from the platform. If an attacker can tap into these side channel leakages, they will gain access to a large volume of data to aid in their venture. It is suspected that any cryptographic algorithm can be broken with enough traces of its operation [MOP08].

The power [KJJ99] and electromagnetic [Agr+03] fingerprints of a circuit are commonly used to perform SCAs on its cryptographic algorithms. These magnitudes fluctuate quickly enough to give a good indicator of the status of the device. This characteristic, however, implies that sophisticated equipment may be needed to capture the information. After all, this is a typical problem of sampling [PM96]. Quick variations of small circuits may only be captured by expensive oscilloscopes and spectrum analyzers. Which difficulties the task of the attacker. They not only require direct access to the device; they must also count on the necessary equipment to mount an attack. Or so we thought. Recent works have shown that attackers do not require direct access to the target platform [ZS18; Sch+21] nor specialized analysis devices [Gra+19; Gra+21] to mount power analysis attacks on internet-enabled platforms. If we remember that a current trend of technology is their hyper-connectivity [CB21], suddenly these devastating attacks become more concerning.

Cryptographic algorithms can be protected against SCAs by employing different obfuscation techniques such as masking [PR13] and shuffling [Vey+12].

These countermeasures, however, are adjusted for every algorithm and require extensive testing to confirm their effectiveness [Bay+11]. Additionally, some mitigation strategies may be favored for some platforms in function of the system constraints. This implies that multiple solutions may need to be evaluated to come up with the best choice. Developing and implementing multiple prototypes can burden any project with significant expenses. Computer aided design may supply a solution to some of these problems. This strategy could be adopted to perform many tests with small implementation costs.

The *gem5* simulator [Bin+11] can emulate the operation of multiple processors including ARM and RISC-V. This is done not only in a behavioral sense but modeling the platforms at the level of micro-architectural components. *gem5* has been used to emulate attacks against micro-architectures such as RowHammer [Fra+21] and Spectre [AM21]. The main interest for studying these attacks in *gem5* is that this software is open source and can be easily customized to model different hardware components. For example, it is possible to emulate custom hardware accelerators along complex processor systems. It is also possible to parameterize the simulation and customize the system to test multiple architectures. Therefore, it is a useful tool to test multiple target architectures and implement and test potential countermeasures.

A *gem5* simulation creates a large set of statistics associated with the operation of the system. These data can be used to study the behavior of the different components in the simulated platform. For example, we can estimate the instructions processed by the ALU and the memory access activity. These estimations will depend on the data being processed by the core or even on the activity of any custom accelerators that can perform direct-memory access. In our case, we focus on the first scenario. As the activity of the multiple underlying components will affect the power consumption of the entire system, we propose that the simulation statistics can be used to approach the electrical behavior of a simulated platform. Therefore, if there is any correlation between the metrics of the simulation and the application data, then *gem5* could be employed to emulate SCAs.

In this paper, we use *gem5* to perform the emulation of an ARM processor with a Linux-based operating system. We employ the platform described in [For+21] which features TrustZone [AF04] and OP-TEE [Bec14] as logical protections. Using this system, we implement different cryptographic operations and analyze the process for retrieving and processing the simulation statistics. With these data we try to perform correlation power analysis, simple power analysis, and power-based covert channels. To the best of our knowledge, this is the first time such an analysis has been conducted. Our findings show that the data acquisition process is particularly challenging to conduct some of these evaluations. Nonetheless, we show points of utility for the improvement of the simulator to make the proposed approach more practical.

The rest of the paper is structured as follows. In Section 2 we discuss related works from literature. In Section 3 we describe our methods and the results obtained. Lastly, Section 4 presents a discussion of our findings and concludes this work. As an appendix, Section 5 has further details on our experimental work.

2 State of the Art

The literature reports multiple analysis tools which allow to study the energy footprint of a circuit before it is implemented. They intend to model the power dissipation of the design based on different approaches. These power estimations can be used, in some cases by the same tool, to detect information leakage in the design.

Some of the most used power analysis solutions have been created by technology vendors and thus are specialized for their products. We can mention, for example, *Incisive Palladium III* by *Cadence Design Systems* and *XPower Analyzer* by *AMD-Xilinx*. The former allows to study the dynamic power dissipation of ASIC designs and the latter supplies static and dynamic power estimation for FPGAs. *Incisive Palladium III* employs the gate model of the chip whereas *XPower Analyzer* employs the RTL model of the design. Therefore, it is necessary to design and describe the architecture to use these tools.

Other tools which employ the gate model of the circuit to study potential leakage sources include PARAM [Ars+20], ACA [Yao+20], and CASCADE [Šij+20]. These tools analyze the signals or gates in the design to estimate the points of interest for side channel analysis. The main drawback for ACA and PARAM is that they are closed-source projects, and while CASCADE claims to be open source it relies on commercial EDA tools such as *PrimeTime* by *Synopsys*. RTL-PSC [He+19] is a tool similar in functionality, however it only requires the RTL specification of the circuit. This system also relies on proprietary software, in particular *VCS* by *Synopsys*.

Leakage verification is another approach for evaluating the vulnerability of a circuit against SCAs. It relies on formal methods for assessing the information leakage such as Hamming models and TVLA. Several tools in the market employ the gate model of the platform for conducting such analyses. We can mention SCRIPT [Nah+20] and PATCH [SJ22] in this category. Others like AMASIVE [HSZ13] and KARNA [Slp+19] employ higher abstractions and thus require a lower design investment for their use. All these systems are proprietary.

Open-source tools for leakage verification include MAPS [LGD18] and COCO [Gig+21]. The former is exclusive for ARM Cortex M3 systems but is not cycle accurate. The latter can analyze any circuit but requires a gate-level description of the platform. Another tool which relies on an open-source initiative is SLEAK [WHK14]. This system employs *gem5* to perform the emulation of ARM Cortex A8 processors. However, SLEAK itself is not readily available.

The general problem with these tools is that most of them are in some part closed-source projects. Furthermore, most of them are meant to study simple models of the system under test and not complete architectures. *gem5*, on the other hand, can emulate large architectures with relative ease. Another drawback of most of the tools in the literature is that they perform leakage analysis or verification over the gate model of the system under study. This means that it is necessary to design and describe the platform to later analyze it. However, this process implies a significant investment which we intend to prevent. In contrast *gem5* only requires a functional model of the architecture thus effectively reducing the design cost.

3 Methods

In this Section we describe the virtual architecture simulated in *gem5*. We illustrate our method for retrieving the statistics of operation and how to process the information. Finally, we describe our experiments in analyzing these data to emulate SCAs.

3.1 The simulated platform

This work uses the virtual and open platform from [For+21] that simulates the behavior of micro-architectural features and their interactions with the peripherals, like accelerators and memories in emerging technologies.

A *gem5* simulation is composed of Python configuration files and C++ program files, which are compiled into a *gem5* binary. The simulator behaves like a Python interpreter. The configuration files connect Python objects, which are either other Python objects or *gem5* primitives standing for hardware modules, through ports generally using a *gem5* memory packet protocol. The hardware primitives can be configured using parameters directly in the Python configuration file. We can, for example, specify the type of processors to be used, the size and technology of the cache, and even include *ad-hoc* components to stand for hardware accelerators.

One of the interesting characteristics of *gem5* is that it allows to model different processor architectures (x86, ARM, RISC-V) and specify multiple optimization levels (`debug`, `opt`, `fast`) to study multiple characteristics of the system. In our work, we focus on ARM platforms and use the `opt` level to conduct our experiments. Another detail to consider is that the simulator can perform either *system-call emulation* or *full simulation* of a platform. We have chosen the latter as this method allows us to obtain statistics which are closer to those found in a physical device.

We have simulated a single-core system with an HPI processor model, which extends the conventional `MinorCPU` (a *timing* model). We chose to use a single core to reduce the complexity of the simulation and the number of statistics produced. The system was provided with L1 and L2 caches as well as a memory management unit (MMU). We added a 2GB DDR3_1600_8x8 unit which was emulated using Ramulator [KYM15]. We employed the `VExpress_gem5_Foundation` machine type which allows to execute the `ArmTrustedFirmware` workload [Her20]. As in [For+21] we have also included the OP-TEE runtime in the simulation. This allowed us to study the behavior of a *protected* system and analyze its resilience against power attacks. Figure 1 illustrates the characteristics of the simulated platform.

From every execution of the simulation, *gem5* dumps statistics like the number of `IntAlu` instructions executed or the frequency of the processor, by default with a rate of $1E - 3$. That is, every simulated millisecond a statistics file is updated with different measurements. This *sampling* rate can also be adjusted in the configuration file for the simulation. In this way it is possible to generate data at a rate consistent with the sampling theorem [PM96]. For example, if the core is running at 1MHz then it will be necessary to use a dump rate under $5E - 7$.

The main problem with this approach is that bumping the dump rate increases the complexity of the simulation and the volume of data produced. On




Figure 1: The architecture under analysis.


the other hand, reducing the frequency of the core causes the simulation to behave erratically. The lowest frequency we managed to employ was 2MHz, but it was impossible to perform any significant computations with the system. With our evaluation platform (Ubuntu 20.04 LTS, 11th Gen Intel Core i7 @ 3.00GHz \times 8, 32 GiB RAM, *gem5* v21) we found that it was possible to use a core frequency of 10MHz and an acquisition rate of 25MHz. Even then, some statistics files reached the size of 1TB. We developed basic Python and Matlab scripts to separate all the statistics into independent time-series and analyze the data.

The main idea of using the simulation statistics to approximate the power dissipation of the system comes from [Han22]. A simple non-representative power model can be approximated as shown in Equation 1. Evidently, the dynamic component of the model would be the most useful for power analysis. In Figure 2, we illustrate how the statistics retrieved from a *gem5* simulation can be used to approach a power estimation.


$$\begin{aligned}
 pow &= dyn + stat \\
 dyn &= V \times (2A \times ipc + 3pA \times dentry_misses) \\
 stat &= 4 \times temperature
 \end{aligned} \tag{1}$$

An important challenge in the analysis of the samples obtained from the operation of a processor is to accurately find the different *parts* of its program. As illustrated in Figure 2b, we overcome this drawback by changing the core frequency at precise points of its operation. This is performed with the same mechanism we use to decrease the operational frequency to follow the sampling theorem.

Another problem with this approach is that the goodness-of-fit for a power estimation will depend greatly on the technology of the system, the architecture, and the mathematical model employed. Creating a function which can accurately describe the behavior of a physical device is equivalent to the design of its digital twin. Which is a problem outside of the scope of this work. However, as shown in Equation 1, even the most precise power model must be created from a composition of readily available statistics. If there is information concerning the operation of the system in these data, then this information would also be embedded in the power estimation. So, investigating the use of the raw statistics for emulating power attacks can be considered as an equivalent problem.



(a) *gem5* statistics



(b) estimation of dynamic power

Figure 2: Statistics obtained from a full-system *gem5* simulation of the architecture illustrated in Figure 1 with a sampling period of $1E-5$. The simulation included the boot sequence, the initialization of OP-TEE and the execution of `optee_example_aes`. The vertical axis shows the normalized amplitude since all the statistics have dynamic ranges which differ significantly; we simply normalize each trace for display purposes.

3.2 Power dissipation analysis

When we transition from the analysis of a power trace to the analysis of the statistics which may influence the power trace, we are confronted with a significant challenge: the volume of the data. In a regular *gem5* simulation we can find over 1,000 different metrics to analyze, which is time consuming. It is necessary to find out which statistics carry the most significant information for performing power analysis.

For this Subsection we used as case study the power analysis of AES [NIS01], since it is a well-known subject in the literature. However, for a preliminary assessment, obtaining enough data from multiple *full* executions of AES seemed impractical. So, we simplified the algorithm under analysis to the essential operations used in SCA attacks. The pseudo-code for our simplified version of AES is provided in Algorithm 1.

Algorithm 1 Algorithm under power analysis

Require: k , an 8-bit random integer

Require: SBOX, the substitution box of AES

Require: f, f' , two frequencies of the core with $f' < 2f_s$

for $i = 0$ **to** 255 **do**

 cpu_freq $\leftarrow f'$ {Pull trigger}

 SBOX($i \oplus k$)

 cpu_freq $\leftarrow f$ {Release trigger}


 wait

end for


We selected two key values (0x2B, 0x7E) to analyze whether it was possible to find useful information in the *gem5* output. For that, we obtained the simulation statistics of the platform illustrated in Figure 1 when executing the program from Algorithm 1. From this, we obtained over 800 files holding multiple metrics from the simulation such as the number of instructions per cycle and the number of cache accesses; with some post-processing (removing incomplete and all-zero files) this number was reduced to 150. Evidently, it was not practical to process all these data with the usual statistical techniques employed for power analysis. So, we devised a strategy to reduce the number of targets for study:

STEP 1: Auto-correlation Compute the correlation between every statistics file within a set (for example, the set of all the statistics obtained from the use of the key 0x2B) to discard those with high similarity. Evidently, if some measurements are similar, they will have a high correlation and provide about the same information. We found 151 stat pairs with a correlation coefficient over 0.95, involving 23 files. We were still left with over 100 files to evaluate. This is shown in Figure 5.1.

STEP 2: Cross-correlation Compute the correlation between every statistics file in a set against the respective statistics in a distinct set (that is, the statistics for the key 0x2B against the statistics for the key 0x7E) as to drop those with high similarity. From this experiment we found only 30 stats with a



(a) sampling period is $2.5E - 7$



(b) sampling period is $4E - 8$

Figure 3: A simple analysis of the clock statistics from multiple experiments illustrates how under regular conditions (Figure 3a) the behavior of *gem5* is consistent. However, when the core frequency is decreased (to 10MHz in this case) the execution of the program becomes irregular (Figure 3b).

correlation coefficient under 0.5 across sets, and only six stats with a correlation coefficient under 0.3 across sets. This is shown in Figure 5.2.

Then, we computed the correlation matrix of the statistics for one of the keys against their own set (auto-correlation) and the correlation matrix between the statistics in both sets (cross-correlation). The first experiment would give us information about the dependence between statistics and the second would highlight which statistics held more information related to the data processing. These findings are detailed in Section 5.

At first sight these results seemed promising as it was possible to find which statistics were more interesting to evaluate. However, in practice, when we set $f' = 10MHz$ and $t_s = 4E - 8$, as to follow the sampling theorem, the fidelity of the analysis decreased. As shown in Figure 3, the behavior of the simulation became inconsistent. This could be further corroborated with an analysis of the cross-correlation between statistics for different key values. For the experiment in Figure 3a the main diagonal of the correlation matrix showed a mean $\delta = 0.72$ and standard deviation $\sigma = 0.24$, while for the experiment in Figure 3b the respective values were $\delta = 0.47$ and $\sigma = 0.28$. Notably, the correlation coefficient between the clock signals decreased from 1.00 to 0.56. Another issue we met was that for the experiment in Figure 3b we could only capture 20 traces. This task took over 72h for each key value and yielded 1TB of statistics in each case.


Despite these findings, we decided to test whether it was possible to perform Test Vector Leakage assessment [SM15] over the data (evaluating the complete sets of stats). This should have helped to find out if the leakages we saw were dependent or not of the secret value. For this experiment, the secret values were the key values `0x2B` and `0x7E` (these values are found in the generic AES key used to obtain test vectors) and we performed a **Student's *t*-test** to see if the leakages' distribution differed from one another. However, using the raw leakages, no statistics overtook the leakage detection threshold of 4.5. This confirmed the null hypothesis of the *t*-test that both sets followed the same distribution. Supposing leakages are independent of the key used, it seems impossible to mount a successful differential attack using single *gem5* statistics. These negative results may change if we can get a larger set of traces for each key, apply signal processing techniques to improve the synchronization of the traces, or consider attacks on various statistics at the same time.

3.3 Timing analysis


Even though we were unable to show a clear relationship between the *gem5* statistics and the secret key of AES, it was clear from Figure 2b that there was some influence of the algorithm in the model. So, we decided to evaluate whether we could perform simple power analysis using the simulator.

It is well known that unprotected implementations of algorithms like RSA [NIS23] and elliptic curve cryptosystems can leak the secret key via a simple timing analysis. This is most commonly due to the use of naive multiplication algorithms which perform branching operations as a function of the key value.

Given the glaring drawbacks of *gem5* for simulating complex operations we decided to skip on RSA. Instead, we analyzed a simple binary-field double-and-add 1024-bit multiplication implemented in C. The algorithm under analysis



(a) target frequency of 600MHz



(b) target frequency of 10MHz

Figure 4: A timing analysis of the execution of the program in Algorithm 2 with a sampling frequency of 25 MHz. The vertical axis shows the normalized amplitude since all the statistics have dynamic ranges which differ significantly; we simply normalize each trace for display purposes.

for this scenario is provided in Algorithm 2. We conducted two experiments to evaluate our hypothesis, see Figure 4.

Algorithm 2 Algorithm under timing analysis

Require: $g(t)$, an irreducible polynomial for \mathbb{K}

Require: $a, b \in \mathbb{K}$, two random integers

Require: f, f' , two frequencies of the core with $f' < 2f_s$

cpu_freq $\leftarrow f'$

for $i = 0$ **to** 5 **do**

$a \times b \in \mathbb{K}$

end for

cpu_freq $\leftarrow f$


In the first experiment (Figure 4a) we set the core frequency to 600 MHz and used an acquisition rate of $4E-8$. This allowed us to see the activity of the core during processing. However, we could not find any timing pattern which corresponded with the execution of the field multiplications. So, we decreased the core frequency to investigate whether following the sampling theorem was necessary to find any timing behavior. But the result was just the opposite. As illustrated in Figure 4b the statistics produced were of inferior quality. For example, in the case of the IPCs, the number of NAN values increased from less than 1% to over 60%. Once again, we arrived at the conclusion that *gem5* simulations behave inconsistently when low frequencies are used.

3.4 Covert channel attacks


SCAs do not only focus on retrieving side channel information. Another vulnerability reported in the literature [Gna+21] is to use the power footprint of the platform to transfer information covertly. In this attack model, the adversary leverages the use of components in the platform to produce a discernible pattern in the energy consumption of the circuit. For example, hardware components, accelerators, and memory elements can be activated or accessed with a pattern to transfer a message. The receiver is generally another component within the system which would not normally be allowed to share a communications channel with the sender. In the presence of trusted execution environments like OP-TEE, the exclusion level of different peripherals or accelerators can mitigate such attacks.

However, to which extent different components can affect the power dissipation of the device is unknown until tested. As we illustrated in Figure 2b, even the execution of cryptographic components can produce a noticeable impact on the statistics of the simulation. We conducted an experiment to illustrate the issue, see Figure 5.

In this scenario we simply executed the `optee_example_aes` with a basic “test message” input once or twice, repeatedly. Each instance of the example performed encryption and decryption requests to the AES TA to encrypt the data and verify if the computation was correct. As can be seen from Figure 5a, the operations performed by the AES TA have a significant impact on the statistics reported by the simulator, even with a low sampling frequency. The activation pattern is reflected even in simple power models like the one included



(a) *gem5* statistics



(b) estimation of dynamic power

Figure 5: Statistics obtained from the execution of `optee_example_aes` with a sampling period of $1E - 5$. The vertical axis shows the normalized amplitude since all the statistics have dynamic ranges which differ significantly; we simply normalize each trace for display purposes.

in Figure 5b. So, we expect that using the OP-TEE trusted applications as power wasters would have a similar effect on a physical device.

It is interesting to note that for this experiment we didn't request any frequency changes to the kernel. The frequency variation seen in Figure 5 is the result of the decisions taken by the `schedutil` governor of Linux. But despite this interference the pattern produced to encode a secret message is clearly distinguishable and could be easily decoded with the application of some filters and basic heuristics.

4 Conclusions

In this paper we have investigated the feasibility of using the *gem5* simulator to emulate SCAs on microprocessor systems. These devices are commonly found in a wide range of platforms which makes them an interesting case study. In recent years, their security resilience against physical attacks has been scrutinized as this kind of vulnerabilities has shown prowess for compromising the security of such systems. It has been proven that SCAs need more study with the emergence of novel attack models like remote power analysis and covert channels.

The use of a simulator would allow us to solve important challenges found in the study of power attacks. It would allow the security auditors to review a large combination of computing platforms, architectures, and security algorithms. All without the need to invest time and resources in the creation of prototypes and the implementation of the solutions. *gem5* has the potential to fill this niche as a simulator which allows to study multiple aspects of the platform. It can be used to derive precise statistics about the operation of the core and then combine them to obtain more complex estimators for the system. However, some challenges must be overcome to improve the usability of this platform:

- it is necessary to improve the production of statistics to make their reporting more concise
- it is necessary to implement mechanisms for enabling or disabling statistics with fine granularity

While the current distribution of the simulator includes some rudimentary support to achieve these goals, this code could very well be considered legacy. Its use is obscure, not documented, and requires us to manually edit core scripts of the simulator. This is followed by a recompilation of the software, which is a lengthy process. Despite these issues, we investigated the possibility of emulating power SCAs using *gem5*.

In relation to power analysis (correlation power analysis, simple power analysis) we showed a relationship between the algorithms and the simulation statistics. This allowed us to find certain statistics of interest. However, upon further testing, we could not verify that they held a statistical relationship with the secret key of the algorithms under study. The main challenge for conducting these experiments was the simulation time and the data volume. As a possible solution we explored the idea of decreasing the operational frequency of the virtual platform, but we found out that *gem5* has problems emulating low-frequency systems. Alternatively, if it was possible to select which statistics are produced it would not be necessary to decrease the target frequency. Reducing

the number of statistics reported would also contribute to reducing the simulation time. On the other hand, our study on power-based covert channels clearly showed the potential of the simulator for studying architectural and software vulnerabilities.

Data availability

The multiple sets of statistics used in our experiments as well as the scripts created for processing the data can be accessed freely on <https://github.com/CarlosAndresLARA/power-gem5>.

Acknowledgements

This work has been supported by the French government through the *Agence Nationale de la Recherche* under project ARCHISEC (ANR-19-CE39-0008) and in the framework of the *France 2030* initiative under project ARSENE (ANR-22-PECY-0004).

References

- [AF04] Tiago Alves and Don Felton. “Trustzone: Integrated hardware and software security”. In: *Information Quarterly* 3.4 (2004), pp. 18–24.
- [Agr+03] Dakshi Agrawal, Bruce Archambeault, Josyula R. Rao and Pankaj Rohatgi. “The EM Side—Channel(s)”. In: *4th International Workshop on Cryptographic Hardware and Embedded Systems (CHES)*. Springer, Aug. 2003, pp. 29–45. DOI: 10.1007/3-540-36400-5_4.
- [AM21] Pierre Ayoub and Clémentine Maurice. “Reproducing Spectre Attack with Gem5: How To Do It Right?” In: *14th European Workshop on Systems Security*. Association for Computing Machinery, Apr. 2021, pp. 15–20. DOI: 10.1145/3447852.3458715.
- [Ars+20] Muhammad Arsath K F, Vinod Ganesan, Rahul Bodduna and Chester Rebeiro. “PARAM: A Microprocessor Hardened for Power Side-Channel Attack Resistance”. In: *2020 IEEE International Symposium on Hardware Oriented Security and Trust*. Dec. 2020, pp. 23–34. DOI: 10.1109/HOST45689.2020.9300263.
- [Bay+11] Ali Galip Bayrak, Francesco Regazzoni, Philip Brisk, François-Xavier Standaert and Paolo Ienne. “A First Step towards Automatic Application of Power Analysis Countermeasures”. In: *48th Design Automation Conference*. ACM, June 2011, pp. 230–235. DOI: 10.1145/2024724.2024778.
- [Bec14] Joakim Bech. *How to create and run Trusted Applications on OP-TEE*. [Online] <https://www.slideshare.net/linaroorg/lcu14103-how-to-create-and-run-trusted-applications-on-optee>. Jan. 2014. (Visited on 11/15/2022).

- [Ber05] Daniel J. Bernstein. *Understanding brute force*. ECRYPT STVL Workshop on Symmetric Key Encryption. Apr. 2005. URL: <https://cr.yp.to/snuffle/bruteforce-20050425.pdf>.
- [Bin+11] Nathan Binkert et al. “The Gem5 Simulator”. In: *ACM SIGARCH Computer Architecture News* 39.2 (Aug. 2011), pp. 1–7. DOI: 10.1145/2024716.2024718.
- [BR05] Mihir Bellare and Phillip Rogaway. *Introduction to Modern Cryptography*. University of California. [Online] <https://web.cs.ucdavis.edu/~rogaway/classes/227/spring05/book/main.pdf>. May 2005. URL: <https://web.cs.ucdavis.edu/~rogaway/classes/227/spring05/book/main.pdf>.
- [CB21] Prasenjit Chanak and Indrajit Banerjee. “Internet-of-Things-Enabled SmartVillages: An Overview”. In: *IEEE Consumer Electronics Magazine* 10.3 (July 2021), pp. 12–18. DOI: 10.1109/MCE.2020.3013244.
- [For+21] Quentin Forcioli et al. “Virtual Platform to Analyze the Security of a System on Chip at Microarchitectural Level”. In: *5th Workshop on Security of Software/Hardware Interfaces*. IEEE, Sept. 2021, pp. 96–102. DOI: 10.1109/EuroSPW54576.2021.00017.
- [Fra+21] Loïc France, Florent Bruguier, Maria Mushtaq, David Novo and Pascal Benoit. “Implementing Rowhammer Memory Corruption in the gem5 Simulator”. In: *2021 IEEE International Workshop on Rapid System Prototyping*. Oct. 2021, pp. 36–42. DOI: 10.1109/RSP53691.2021.9806242.
- [Gig+21] Barbara Gigerl, Vedad Hadzic, Robert Primas, Stefan Mangard and Roderick Bloem. “Coco: Co-Design and Co-Verification of Masked Software Implementations on CPUs”. In: *30th USENIX Security Symposium*. USENIX Association, Aug. 2021, pp. 1469–1468.
- [Gna+21] Dennis R. E. Gnad, Cong Dang Khoa Nguyen, Syed Hashim Gillani and Mehdi B. Tahoori. “Voltage-Based Covert Channels Using FPGAs”. In: *ACM Transactions on Design Automation of Electronic Systems* 26.6 (June 2021), pp. 1–25. DOI: 10.1145/3460229.
- [Gra+19] Joseph Gravellier, Jean-Max Dutertre, Yannick Teglia and Philippe Loubet-Moundi. “High-Speed Ring Oscillator based Sensors for Remote Side-Channel Attacks on FPGAs”. In: *2019 International Conference on ReConFigurable Computing and FPGAs (ReConFig)*. IEEE, Dec. 2019, pp. 1–8. DOI: 10.1109/ReConFig48160.2019.8994789.
- [Gra+21] Joseph Gravellier, Jean-Max Dutertre, Yannick Teglia and Philippe Loubet Moundi. “SideLine: How Delay-Lines (May) Leak Secrets from Your SoC”. In: *2021 International Workshop on Constructive Side-Channel Analysis and Secure Design*. Springer, Oct. 2021, pp. 3–30. DOI: 10.1007/978-3-030-89915-8_1.
- [Han22] Thomas E. Hansen. *ARM Power Modelling*. gem5 Documentation. [Online] https://www.gem5.org/documentation/learning_gem5/part2/arm_power_modelling/. Nov. 2022. URL: https://www.gem5.org/documentation/learning_gem5/part2/arm_power_modelling/.

- [He+19] Miao He et al. “RTL-PSC: Automated Power Side-Channel Leakage Assessment at Register-Transfer Level”. In: *IEEE 37th VLSI Test Symposium*. July 2019, pp. 1–6. DOI: 10.1109/VTS.2019.8758600.
- [Her20] Adrian Herrera. *Running Trusted Firmware-A on gem5*. arm Community. [Online] <https://community.arm.com/arm-research/b/articles/posts/running-trusted-firmware-a-on-gem5>. June 2020. URL: <https://community.arm.com/arm-research/b/articles/posts/running-trusted-firmware-a-on-gem5>.
- [HSZ13] Sorin A. Huss, Marc Stöttinger and Michael Zohner. “Amasive: an adaptable and modular autonomous side-channel vulnerability evaluation framework”. In: *Number Theory and Cryptography* (2013), pp. 151–165.
- [KJJ99] Paul Kocher, Joshua Jaffe and Benjamin Jun. “Differential Power Analysis”. In: *19th Annual International Cryptology Conference (CRYPTO)*. Springer, Aug. 1999, pp. 388–397. DOI: 10.1007/3-540-48405-1_25.
- [Kru12] Robert Krulwich. *Which Is Greater, The Number Of Sand Grains On Earth Or Stars In The Sky?* NPR. [Online] <https://www.npr.org/sections/krulwich/2012/09/17/161096233>. Sept. 2012. URL: <https://www.npr.org/sections/krulwich/2012/09/17/161096233/which-is-greater-the-number-of-sand-grains-on-earth-or-stars-in-the-sky>.
- [KYM15] Yoongu Kim, Weikun Yang and Onur Mutlu. “Ramulator: A fast and extensible DRAM simulator”. In: *IEEE Computer Architecture Letters* 15.1 (Mar. 2015), pp. 45–49. DOI: 10.1109/LCA.2015.2414456.
- [LGD18] Yann Le Corre, Johann Großschädl and Daniel Dinu. “Micro-architectural power simulator for leakage assessment of cryptographic software on ARM Cortex-M3 processors”. In: *9th International Workshop Constructive Side-Channel Analysis and Secure Design (COSADE)*. Springer. Apr. 2018, pp. 82–98.
- [MOP08] Stefan Mangard, Elisabeth Oswald and Thomas Popp. *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. Advances in information security. Springer, 2008. ISBN: 978-0-387-38162-6.
- [Nah+20] Adib Nahiyan et al. “SCRIPT: A CAD Framework for Power Side-Channel Vulnerability Assessment Using Information Flow Tracking and Pattern Generation”. In: *ACM Transactions on Design Automation of Electronic Systems* 25.3 (May 2020). ISSN: 1084-4309. DOI: 10.1145/3383445.
- [NIS01] NIST. *FIPS 197: Advanced Encryption Standard (AES)*. Standard. [Online] <https://csrc.nist.gov/publications/detail/fips/197/final>. National Institute of Standards and Technology, Nov. 2001. DOI: 10.6028/NIST.FIPS.197-upd1.ipd.
- [NIS23] NIST. *FIPS 186-5: Digital Signature Standard (DSS)*. Standard. [Online] <https://csrc.nist.gov/publications/detail/fips/186/5/final>. National Institute of Standards and Technology, Feb. 2023. DOI: 10.6028/NIST.FIPS.186-5.

- [PM96] John G. Proakis and Dimitris G. Manolakis. *Digital Signal Processing: Principles, Algorithms, and Applications*. Prentice Hall, 1996. ISBN: 978-0-133-94289-7.
- [PR13] Emmanuel Prouff and Matthieu Rivain. “Masking against Side-Channel Attacks: A Formal Security Proof”. In: *32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*. Springer, May 2013, pp. 142–159. DOI: 10.1007/978-3-642-38348-9_9.
- [Sch+21] Falk Schellenberg, Dennis R. E. Gnad, Amir Moradi and Mehdi B. Tahoori. “An Inside Job: Remote Power Analysis Attacks on FPGAs”. In: *IEEE Design & Test* 38.3 (Mar. 2021), pp. 58–66. DOI: 10.1109/MDAT.2021.3063306.
- [Šij+20] Danilo Šijačić, Josep Balasch, Bohan Yang, Santosh Ghosh and Ingrid Verbauwhede. “Towards efficient and automated side-channel evaluations at design time”. In: *Journal of Cryptographic Engineering* 10 (June 2020), pp. 305–319.
- [SJ22] Vahhab Samadi Bokharaie and Ali Jahanian. “Power side-channel leakage assessment and locating the exact sources of leakage at the early stages of ASIC design process”. In: *The Journal of Supercomputing* (June 2022), pp. 1–26.
- [Slp+19] Patanjali Slpsk, Prasanna Karthik Vairam, Chester Rebeiro and V. Kamakoti. “Karna: A Gate-Sizing based Security Aware EDA Flow for Improved Power Side-Channel Attack Protection”. In: *2019 IEEE/ACM International Conference on Computer-Aided Design*. Nov. 2019, pp. 1–8. DOI: 10.1109/ICCAD45719.2019.8942173.
- [SM15] Tobias Schneider and Amir Moradi. “Leakage Assessment Methodology”. In: *17th International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, Sept. 2015, pp. 495–513. ISBN: 978-3-662-48324-4.
- [Vey+12] Nicolas Veyrat-Charvillon, Marcel Medwed, Stéphanie Kerckhof and François-Xavier Standaert. “Shuffling against Side-Channel Attacks: A Comprehensive Study with Cautionary Note”. In: *18th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*. Springer, Dec. 2012, pp. 740–757. DOI: 10.1007/978-3-642-34961-4_44.
- [WHK14] Dan Walters, Andrew Hagen and Eric Kedaigle. *SLEAK: A side-channel leakage evaluator and analysis kit*. Tech. rep. AD1107774. MITRE CORP BEDFORD MA, Jan. 2014.
- [Yao+20] Yuan Yao, Tarun Kathuria, Baris Ege and Patrick Schaumont. “Architecture Correlation Analysis (ACA): Identifying the Source of Side-channel Leakage at Gate-level”. In: *2020 IEEE International Symposium on Hardware Oriented Security and Trust*. Dec. 2020, pp. 188–196. DOI: 10.1109/HOST45689.2020.9300271.
- [ZS18] Mark Zhao and G. Edward Suh. “FPGA-Based Remote Power Side-Channel Attacks”. In: *2018 IEEE Symposium on Security and Privacy*. IEEE, May 2018, pp. 229–244. DOI: 10.1109/SP.2018.00049.

5 Correlation analysis for the statistics

We performed a correlation analysis for a reduced set of statistics (sorted alphabetically). From the autocorrelation shown in Figure 5.1 we could figure out (by looking outside the main diagonal) that there seem to be some statistics with high similarity (*SimdAdd*, *SimdAlu*, *SimdCmp*) which may carry redundant information.




Figure 5.1: Correlation within the set of statistics generated from the usage of the key 0x2B.

From the cross-correlation shown in Figure 5.2 we can figure out (also

by looking at the main diagonal) that there are statistics with low similarity (*FloatMisc*, *No_OpClass*) which may be influenced by data dependency. In the cross-correlation matrix we can observe a slight asymmetry resulting from the analysis of equivalent pairs produced by taking one statistic from each set (e.g. $\text{corr}(\text{set}_A.\text{stat}_A, \text{set}_B.\text{stat}_B)$ vs $\text{corr}(\text{set}_B.\text{stat}_A, \text{set}_A.\text{stat}_B)$).




Figure 5.2: Correlation between two sets of statistics generated from the usage of the keys 0x2B and 0x7E.

It is also interesting to note how the *clock* statistic (which carries the *alignment* information) has a high cross-correlation value, which lends validity to the proposed method for segmenting the traces.