

# Scrooge Attack: Undervolting ARM Processors for Profit

(Practical experience report)

Christian Göttel\*, Konstantinos Parasyris<sup>†</sup>, Osman Unsal<sup>‡</sup>, Pascal Felber\*, Marcelo Pasin\*, Valerio Schiavoni\*

<sup>†</sup>Lawrence Livermore National Laboratory, parasyris1@llnl.gov

<sup>‡</sup>Barcelona Supercomputing Center, osman.unsal@bsc.es

\*Université de Neuchâtel, Switzerland, first.last@unine.ch

**Abstract**—Latest ARM processors are approaching the computational power of x86 architectures while consuming much less energy. Consequently, supply follows demand with Amazon EC2, Equinix Metal and Microsoft Azure offering ARM-based instances, while Oracle Cloud Infrastructure is about to add such support. We expect this trend to continue, with an increasing number of cloud providers offering ARM-based cloud instances.

ARM processors are more energy-efficient leading to substantial electricity savings for cloud providers. However, a malicious cloud provider could intentionally reduce the CPU voltage to further lower its costs. Running applications malfunction when the undervolting goes below critical thresholds. By avoiding critical voltage regions, a cloud provider can run undervolted instances in a stealthy manner.

This practical experience report describes a novel attack scenario: an attack launched by the cloud provider against its users to aggressively reduce the processor voltage for saving energy to the last penny. We call it the Scrooge Attack and show how it could be executed using ARM-based computing instances. We mimic ARM-based cloud instances by deploying our own ARM-based devices using different generations of Raspberry Pi. Using realistic and synthetic workloads, we demonstrate to which degree of aggressiveness the attack is relevant. The attack is unnoticeable by our detection method up to an offset of  $-50$  mV. We show that the attack may even remain completely stealthy for certain workloads. Finally, we propose a set of client-based detection methods that can identify undervolted instances. We support experimental reproducibility and provide instructions to reproduce our results.

**Index Terms**—ARM, undervolting, attack, detection

## I. INTRODUCTION

Cloud providers continuously upgrade their commercial offerings to adapt to market and customer needs. While the vast majority of them offer computing instances based on x86 processors, the availability of ARM-based cloud instances is quickly expanding. ARM processors are increasing their market share of server-grade machines [1, 2, 3, 4, 5, 6, 7, 8], thanks to additional energy and performance improvements. Amazon [5] deploys ARM-based processors currently shipped in off-the-shelf ARM hardware (their AWS Graviton is essentially a more powerful quad-Raspberry Pi 4B [9]). Scaleway offered instances based on custom-made ARM SoCs with servers smaller than a business card [10]. Table I summarizes a subset of available server-grade ARM processors, supported instruction set architectures (ISA), and providers deploying this hardware. Several generations of ARM processors [1,

Table I: List of server-grade and mimicking ARM processors with their supported ISA. “\*”: used in our evaluation (see §V).

Processor	ISA	Cloud provider
Ampere Altra	ARMv8.2+	Equinix, Oracle
Ampere eMAG 8180	ARMv8	Equinix
AWS Graviton	ARMv8	AWS
AWS Graviton 2	ARMv8.2	AWS
Fujitsu A64FX	ARMv8.2	-
Huawei Kunpeng 920	ARMv8.2	-
Marvell ThunderX	ARMv8	Equinix
Marvell ThunderX2	ARMv8.1	Microsoft Azure
NVIDIA Grace	TBA	-
Broadcom BCM2837(B0)*	ARMv8	-
Broadcom BCM2711*	ARMv8	-

2, 5, 6, 7] are currently available across cloud providers. ARM processors also started reaching into the supercomputing market segment. We expect an increasing availability of ARM Neoverse [11] processors and future server-grade ARM instances to close the performance gap to x86. Processors offer different power management mechanisms to adjust frequencies and voltages. The energy footprint of a single execution step (*i.e.*, one single instruction on a processor) is fairly independent of the CPU frequency but dependent on the CPU voltage [12]. Decreasing the CPU voltage below the nominal value to conserve power is called undervolting<sup>1</sup>. Besides energy savings, undervolting directly influences core temperature and can also reduce core aging [13]. Undervolting, however, incurs the risk of introducing soft [14] and hard-errors related to timing violations [15]. These types of errors can be mitigated by carefully analyzing the guardband of processors [16]. In this practical experience report, we consider a scenario where processors supporting a cloud infrastructure are undervolted by an excessively economic and malicious cloud provider (a *scrooge* §III-A) to profit from additional electricity bill savings, while cloud users (from here on referred to as users) observe similar performance. Unfortunately, undervolting cannot be applied arbitrarily. In fact, it comes at the cost of processor reliability when the supplied voltage is insufficient to drive the processor’s frequency. We believe this is a risk that malicious cloud providers are willing to take. For users, undervolting opens up a new attack vector against their cloud

<sup>1</sup>Notice that Dynamic Voltage and Frequency Scaling (DVFS) differs from undervolting by decreasing frequency as well as voltage.

applications (see our threat model in §III). The main research questions we address in this work are:

**RQ1:** *What is necessary for a malicious cloud provider in order to pull off a stealthy undervolting strategy?*

**RQ2:** *Does a cloud user have the ability to uncover such an undervolting strategy?*

To answer those questions, we need to lay the foundation to better understand consequences of (arbitrary) undervolting, both from the cloud provider and client perspective. In fact, depending on supply voltage, frequency, load, and temperature of the CPU, execution steps can yield erroneous computations. While recent attacks [17, 18] have demonstrated how undervolting can be effectively exploited to gain access to sensitive information, we deal with a different threat model: the infrastructure is undervolted on purpose by a powerful attacker (*i.e.*, the cloud provider), at the risk of exposing hard-to-detect unreliable computing instances for users. Without physical access to instances, nor being able to directly manipulate the supply voltage or frequency, a user’s options remain limited. Nevertheless, a user can adjust the processor’s load and operating performance points (§II-B) to influence its heat dissipation. In order to operate under full load, the processor has to be set to the highest operating performance point, which implies the highest frequency and supply voltage setting. Consequently, undervolted processors present higher probability for erroneous computations to occur because they are unable to maintain high frequencies. This probability is further increased by the propagation delay due to high operating temperature. If erroneous computations result in faults, one can observe application crashes, or kernel panics, leading to cloud instance unavailability. While service level agreements (SLA) [19] typically cover such scenarios, a malicious provider might try to balance its actions to only yield erroneous computations not resulting in faults, basically overcoming SLA protections. For this reason, we designed a non-selective fault injection method for detecting the scrooge attack. The sole purpose of the detection method is to yield intentional application crashes or kernel panics on undervolted instances such that the user is covered by the SLA. While interesting, we consider cloud providers or users exploiting undervolting to leak sensitive information [20, 21] to be out of scope of this work.

Interestingly, ARM-based Raspberry Pis have already been collocated in cloud data centers [22]. With the intent to reproduce and study the dynamics of such deployments (and, to a smaller scale, mimic AWS using ARM nodes), we first study the effects of undervolting on three different ARM processors, focusing on energy savings. Figure 1 shows different normalized energy to throughput ratios (ETR) [12] obtained with ARM Cortex-A processors for the three latest Raspberry Pi models (3B, 3B+, and 4B [9]) at their lowest operational undervolting setting ( $-75$  mV for 3B and 3B+, and  $-15$  mV for 4B) compared to nominal voltage (*i.e.*, 0 mV, no undervolting). As shown, undervolting directly influences energy spent per operation, without negatively affecting throughput. Lower




Figure 1: Normalized energy to throughput ratio (ETR) for undervolted Raspberry Pi model B platforms operating at maximum throughput

normalized ETR values indicate higher energy efficiency for a given throughput. On average across different throughput values we achieved by undervolting 5 % to 13 % better energy efficiency on the 3B and 3B+ and 0 % to 3 % on the 4B. In essence, these results suggest that a cloud provider can indeed undervolt ARM-based instances, without directly compromising the observed performance.

Our contributions are as follows:

- We describe a novel attack scenario based on undervolting by a scrooge cloud provider to lower energy costs.
- We demonstrate how cloud users can with a certain probability detect this novel scrooge attack.
- We provide a temperature-based guardband analysis to narrow down the operation voltage range of an ARM-based processor (§V-D).
- We describe how our analysis can be used to automatically identify undervolted instances (§V-E)
- We present potential energy gains of undervolting systems using a reliability benchmark (§V-F). In general gains can reach up to 37 %.

This practical experience report is organized as follows. Section II provides background on the low-level mechanisms used to undervolt a processor and the Raspberry Pi platform as well as the associated side-effects. Our threat model is given in Section III. We overview our detection method in Section IV. Our in-depth experimental evaluation is presented in Section V. We discuss and review related work in Section VI and Section VII, before concluding in Section VIII.

## II. BACKGROUND

This section defines more precisely a few concepts related to power management (§II-B), *i.e.*, frequency and voltage scaling and associated techniques such as Dynamic Voltage and Frequency Scaling (DVFS) and Adaptive Voltage Scaling (AVS). In §II-D we explain the relation between such techniques and how they affect the overall reliability of a system.

### A. ARM in data centers

Collocation offers allow users to either ship or buy Raspberry Pis in order to deploy lightweight workloads on this




Figure 2: Energy comparison of off-the-shelf and server-grade devices on CPU-bound and memory-bound workloads.

low-energy hardware and thus free up resources on high-energy x86 hardware. Furthermore, Raspberry Pis are the size of credit cards and have much lower cooling demands, which allows hosting a large number of units in a single rack. Such off-the-shelf hardware setups allow for large-scale node deployments as needed in data processing or cloud computing workloads. While off-the-shelf hardware typically lacks in performance and storage capability, its energy consumption remains comparable to server hardware.

Figure 2 compares the energy consumption of ARM-based off-the-shelf hardware (*i.e.*, three different Raspberry Pi models) against server-grade hardware using different x86 architectures. We run a cryptographic (CPU-bound) and a memory allocation (Memory-bound) stressor while measuring the entire device power consumption. The x86 processors used were an AMD EPYC and three different Intel Xeon processor generations, *i.e.*, *Broadwell*, *Kaby Lake* and *Harpertown*. This is a direct comparison of the execution of two distinct binaries of the same source code on two different architectures based on a common metric (J/op). We observe no major difference for CPU-bound operations between different architectures [23]. However, memory-bound operations on off-the-shelf hardware have higher energy consumption. In the case of the Raspberry Pi models, these are due to cache size and memory transfer rate. Nevertheless, off-the-shelf hardware achieves lower energy consumption for both operations compared to older server-grade x86 hardware, *i.e.*, Harpertown. These results indicate that replacing old x86 hardware with recent off-the-shelf ARM-based nodes in data centers will result in energy savings.

### B. Power management

The power dissipated by an integrated circuit depends on static power (leakage current) and dynamic power (switching power). Since about 2005 [24] the power dissipation contribution of dynamic power has become much higher than static power. Nowadays, with the decreased transistor size and lowered threshold voltages, static power is becoming more and more important [25]. In the following, we outline techniques to reduce dynamic power.

**Frequency scaling** regulates (dynamically) the frequency of an integrated circuit in order to change performance, conserve power or reduce the amount of heat dissipation. Reducing the frequency at a constant voltage is called *undervolting* or *throttling*, while increasing the frequency is called *overclocking*. The dynamic power dissipated by an integrated circuit over a period of time is given by  $P = CV^2f$ , where  $C$  is the capacitance,  $V$  is the voltage, and  $f$  is the frequency. Thus, increasing the frequency results in a higher power consumption and operating temperature.

**Voltage scaling** is an open loop system, in which the voltage of an integrated circuit is regulated (dynamically) based on an external setting. Increasing or decreasing the voltage while keeping the frequency constant is called *overvolting* and *undervolting*, respectively. Regulating the voltage enables increasing the frequency or conserving power of an integrated circuit, a particularly useful aspect especially for battery-powered devices. Changing the voltage influences the rate at which capacitances can be charged and discharged. Thus voltage determines the speed and frequency at which an integrated circuit can be operated. Modern operating systems do not provide direct support to adjust a processor’s voltage individually. The processor’s voltage is either regulated by model-specific registers [26] or through firmware.

**DVFS** is the simultaneous software-controlled regulation of voltage and frequency scaling of an integrated circuit. Depending on the process variation (variation of integrated circuits when fabricated) ARM system on a chip (SoC) manufacturers specify a set of operating performance points (OPPs) under worst case conditions. These OPPs are pairs of clock frequencies and voltages under which the integrated circuit is operational with a sufficiently large margin while taking into account thermal conditions. In Linux the CPUFreq kernel driver [27] will chose a set of OPPs based on a specified governor. DVFS has been extensively studied [25, 28, 29] to accelerate multi-threaded applications. x86 manufacturers use their own DVFS implementations [30, 31, 32].

**AVS** [33] is a closed loop system where the voltage is regulated based on its process variation, aging and a feedback loop of sensor data. A hardware monitor or software backed by sensor data determines if the changes made to the system are sufficient or if additional changes are necessary. AVS requires support from both the processor and the power regulators, in order to adjust the voltage accordingly. The Raspberry Pi models B used in this report are equipped with an AVS system.

### C. Raspberry Pi

The Raspberry Pi’s firmware is configured at boot time by a text file containing property-value pairs. For example, the frequency and voltage can be set in this configuration file. A particularity is that voltages can only be set to a nominal offset in steps of 25 mV. This configuration is then parsed by the firmware. This undervolting configuration is specific to the Raspberry Pi and other hardware can more easily be undervolted dynamically at runtime. Notice that the requested CPU frequency in the operating system can deviate from

the actual frequency regulated by the firmware. This is in particular the case if the device reaches the thermal hard limit at 85 °C. Additionally, the 3B+ has a soft limit temperature at 60 °C that will throttle the CPU frequency and voltage.

#### D. Reliability

There are several approaches to determine a processor’s reliability in an undervolted operating regime. Known benchmarks (*e.g.*, SPEC CPU2006 [34], PARSEC [35], *etc.*) are still used [16, 36, 37]. Recently, new specialized software systems [38, 39] have been proposed to maximize power consumption and voltage noise. Even small proof-of-concept programs are sufficient for fault detection under dynamic voltage and/or frequency [17, 20]. Finally, such programs can also be used to characterize the guardband of a system [20, 40].

In this practical experience report we distinguish between three regions with respect to the guardband: safe, critical, and failure. A safe region has a sufficiently high voltage margin, such that erroneous computations or transient faults cannot occur. The critical region designates a small voltage band in which the processor occasionally experiences erroneous results or transient faults. Inside the failure region it is impossible to boot the operating system either because the voltage cannot support the processor’s frequency or because erroneous computations and transient faults lead to kernel crashes or panics.

Undervolted instances become unavailable in case a transient fault leads to an instance crash. From our perspective, current SLAs cover single instances that have crashed because of undervolted hardware, provided users can sufficiently support these claims. The situation is trickier with multiple instances. Deployed instances would have to crash simultaneously, yet process variation plays into the cloud provider’s hands. These crashes are non-deterministic and, therefore, process variation helps obfuscating the undervolted setup. Only simultaneous crashes satisfy today’s cloud provider restrictions in order for users to be covered by the SLA.

### III. THREAT MODEL

In this section we discuss our threat model. In particular, we intend to clarify: (1) which techniques can a malicious cloud provider use to hide an undervolted processor from an unsuspecting user, and (2) which are the methods for a curious user to reveal an undervolted processor? Notice that we validate these methods on a specific hardware configuration (*i.e.*, Raspberry Pi boards using Broadcom BCM2837/BCM2711 processors), but the discussion holds for other platforms relying on similar voltage regulation mechanisms.

#### A. The scrooge cloud provider

We assume the cloud provider has full access to the physical infrastructure and can connect remotely to the physical machines [41]. Furthermore, the cloud provider purposefully undervolts its ARM-based hardware to benefit from additional savings. Firmware configurations can be hidden from users for malicious or security purposes. By maliciously intercepting any voltage reading requests (see §III-C), the cloud provider

ensures that the undervolted state of the cloud infrastructure remains oblivious to users. A cloud provider must find the sweet spot [16] for the undervolt configuration in or near the critical region to provide sufficiently stable instances.

#### B. The curious cloud user

The curious cloud user is suspicious of the cloud provider and intends to uncover its potentially obfuscated activity. Instances of the cloud provider can exclusively be accessed remotely by the user. The only way for a user to detect an undervolted processor is by querying the firmware, normally using a specific executable command file for that. By reading values from the firmware and comparing them to values in the boot configuration file, a user can detect an undervolted processor. If results of firmware queries can be forged, it becomes difficult for a user to uncover the scrooge cloud provider. A confidential and tamper-proof message exchange with the firmware is essential to detect an undervolted processor.

A user can suspect an undervolted processor to operate in the critical region in case of kernel warnings or kernel panics appearing during the system boot or while the system is running, despite these being generic kernel warnings rather than specific ones. In particular if the booting time is longer than expected, then this might hint at a failed boot attempt where the kernel crashed. Most systems have a kernel log that can be consulted by the system administrator. However, a cloud provider can tamper with those kernel logs, and the system utilities are outside the trusted computing base.

#### C. The scrooge attack

The scrooge cloud provider makes undervolted ARM instances available to users. These undervolted instances should be indistinguishable from nominal voltage instances. This includes configuration, firmware, and tools querying CPU voltages. Thus, the undervolt configuration needs to be exchanged for a nominal configuration and any CPU voltage reading request needs to be intercepted. Figure 3 shows different actions the cloud provider has to perform during an instance lifecycle in order to hide the undervolt configuration. When a user boots such an instance, the cloud provider must ensure that the undervolt configuration is loaded by the firmware on the machine the instance is running on. However, this undervolt configuration should not be accessible once the user is connected to the instance. The undervolt configuration has to be swapped for the nominal configuration ❶. Depending on the configuration mechanism, the file system that was booted may be different from the file system the user finds after booting. This includes firmware, operating system kernel, binaries, *etc.* A hidden or obfuscated system service could perform this task while the operating system is booting. An even stealthier approach involves a trusted operating system [42] or auxiliary devices [43] which exchange the configurations before the operating system is booted. Therefore, without proper system attestation, there is no guarantee about the authenticity of the system users believe they have booted. During reboot




Figure 3: State machine with cloud provider actions to obfuscate undervolted machine configuration.

or shutdown of the machine configurations might have to be swapped back ② again.

Any CPU voltage reading request needs to be intercepted and substituted by a plausible nominal voltage value. This will typically involve a kernel driver that will handle the communication with the firmware or accessing model-specific registers. The request can then be intercepted directly in the user space tool or the kernel driver ③. From the kernel driver the request is forwarded ④ and the actual undervolted CPU voltage value is returned to the kernel driver ⑤. The kernel driver then substitutes this value by a some nominal voltage value, *e.g.*, by adding the undervolt offset to the value. The alleged nominal voltage value is then returned to the user ⑥. A more costly but stealthier variant involves the trusted operating system, to which the cloud provider could delegate voltage reading requests instead of a kernel driver.

If users are allowed to deploy their own kernels, then the cloud provider needs a different approach. Voltage reading request can no longer be intercepted in kernel space. Instead, the cloud provider needs to use the hypervisor to intercept CPU voltage requests and substitute them similarly to the kernel driver approach.

In our threat model we assume that the cloud provider will make use of these mechanisms and obfuscate as much as possible the undervolted state of the infrastructure from users, a practical effort with significant benefits. Without access to the firmware configuration nor any untampered message exchange with the CPU voltage regulating mechanism, a user can never be sure to obtain a genuine voltage reading.

#### D. Relevance of discussed techniques

While a scrooge cloud provider has powerful mechanisms in place to hide its undervolted instances, the curious user can still expose this misbehaviour. For instance, the processor’s frequency and package temperature are viable options to test for undervolted conditions. The techniques presented in Section V demonstrate to which extent users can deploy applications stressing aforementioned options on instances and how accurately conclusions can be drawn.

### IV. SCROOGE ATTACK DETECTION

We describe the user’s detection method in this section. Furthermore, we describe under which conditions the detection method works and where difficulties may arise.

We assume that users cannot trust any firmware or system reading on instances. As such, users have no reference to any parameters for adjusting the detection method to the attack.

Users can for this reason make use of simple CPU-bound programs that will put the processor under maximum load while monitoring for faults. Inspired by [17] we propose implementing an arithmetic computation (*i.e.*, multiplication) for which we can validate the result. First we generate two random numbers which are then multiplied until the instance crashes while alternating the position of multiplier and multiplicand. Murdock *et al.* have observed, that the position of the multiplier and the multiplicand can lead to a faulting instruction. After each multiplication the result is compared to the original result. While the processor operates at maximum load it will run at the highest frequency and dissipate heat which will raise its temperature. Under these conditions we achieve the highest probability to inject faults related to timing violations. Depending on the complexity of the RISC circuitry in the ARM processor, certain instructions are more likely to fault than others. To this end, the detection method might not inject faults in its own computation (due to its simple nature), but more likely in other processes. This behavior is favorable, as it allows to run the detection method until the instance itself becomes unavailable due to multiple critical faults in system relevant processes. Thus, we detect an undervolted cloud instance using the detection method by gradually failing processes to crash the instance and make it unavailable.

The detection method depends strongly on how aggressively machines are undervolted and the cooling system employed by the cloud provider. The less a machine is undervolted, the higher the temperature needs to be raised by the detection method to fault processes and *vice versa*. A good cooling system is a lesser problem than a weakly undervolted machine. With a good cooling system the detection method requires a longer time to raise the processor’s temperature. On the one hand implementing a soft limit temperature throttle in order to prevent this detection method is not an ideal solution. Users are less inclined to pay for a service which underperforms compared to alternative services. On the other hand weakly undervolting machines defies the scrooge cloud provider’s original idea of minimizing the electricity bill.

The cloud provider’s options are limited to completely prevent the detection method from unveiling the scrooge attack. Even the powerful setup of the cloud provider to tamper with CPU voltage readings is not sufficient denying the detection method. The scrooge attack has the disadvantage that detection methods have a simple design, but it has the advantage that proving the undervolt state without the firmware is difficult.

### V. EVALUATION

In this section we explore the behavior of Raspberry Pi processors under different nominal and undervolted setups. The information gained from these experiments allows quantifying the attack parameters and determining the type of processes to use for the detection method. Then, we derive the probability at which our detection method can successfully uncover the attack. We begin by describing our experimental setup to undervolt Raspberry Pis before evaluating the firmware’s throttling behavior when reaching the soft limit and limit temperatures.

Table II: Soft limit (SL) firmware throttling on the 3B+

OV level	$V_{\text{arm}}$ [V]	$V_{\text{arm}}^{\text{SL}}$ [V]	$f_{\text{arm}}$ [MHz]	$f_{\text{arm}}^{\text{SL}}$ [MHz]
0	1.3750	1.2688	1400	1200
-1	1.3500	1.2375	1400	1200
-2	1.3188	1.2125	1400	1200
-3	1.2938	1.1875	1400	1200

Table III: Limit temperature (L) throttling on the 3B and 4B

Model	$V_{\text{arm}}^{\text{L}}$ [V]	$f_{\text{arm}}^{\text{L}}$ [MHz]	$f_{\text{core}}^{\text{L}}$ [MHz]
3B	1.2813	{1034, 1087, 1141, 1195, 1200}	{400}
4B	0.8500	{1000, 1500}	{333, 500}

The temperature-based guardband analysis allows detecting the critical region of the device and defines the margin for an undervolt setup. Faults that occurred during the guardband analysis are analyzed to describe the fault injection of the detection method. Finally, we measure the energy efficiency of the undervolted hardware with a reliability benchmark. The dataset gathered for this evaluation will be made publicly available at <https://github.com/ChrisG55/Scrooge-Attack>.

#### A. Experimental settings

We use the three Raspberry Pi models 3B, 3B+, and 4B, while booting from the same SD card a Raspbian Buster distribution (<https://github.com/raspberrypi/linux>). All units rely on recent firmware releases (since June 2020). To simulate a realistic cloud scenario we take all measurements in an air-conditioned room at  $(24 \pm 1)^\circ\text{C}$  and connect the Raspberry Pis to Ethernet and run the SSH daemon. The Raspberry Pis are monitored over UART from an auxiliary machine. No other peripherals are connected to the Pis in order to minimize any interference. Both the Raspberry Pi's and the auxiliary machine's clocks are synchronized using NTP, to easily correlate the power consumption logs recorded on the auxiliary machine to the benchmark running on the Raspberry Pi. The power consumption of the Raspberry Pi is recorded by an Alciom PowerSpy2 [44] over bluetooth.

#### B. Soft limit temperature throttling

We start by evaluating the firmware behavior when reaching the soft limit temperature while running under the CPUFreq *performance* governor. Understanding the throttling behavior helps evaluating the viability of possible mitigation techniques by the cloud provider. The Raspberry Pi documentation mentions that frequency and voltage of the SoC are reduced to decrease heat dissipation but without indicating by how much. The Raspberry Pi 3B+ is the only model with a soft limit temperature programmed into its firmware, therefore, other models are not included in the overvoltage level to OPP mapping reported in Table II. The values indicate that the ARM CPU frequency  $f_{\text{arm}}$  is reduced by 200 MHz and the CPU voltage  $V_{\text{arm}}$  is lowered by about 106 mV (four levels). The voltage stepping of 25 mV remains the same with two exceptions from nominal level  $-1 \rightarrow -2$  with  $-31.2$  mV and from soft limit level  $0 \rightarrow -1$  with  $-31.3$  mV.

#### C. Limit temperature throttling

Next, we evaluate the firmware behavior when reaching the limit temperature while running under the CPUFreq *performance* governor. At the limit temperature, the firmware will throttle the processor to prevent thermal runaway. Notice that model 3B+ is not included here, as it is taking too much time reaching the limit temperature while already being throttled for going beyond the soft limit temperature. Neither the 3B nor the 4B reduce the voltage when reaching the limit temperature as shown in Table III. However, both models reduce their frequency. The 3B is reducing its ARM CPU frequency  $f_{\text{arm}}^{\text{L}}$  in steps of about 54 MHz (except for the first step) while the 4B significantly reduces its frequency by 500 MHz. In addition the 4B also reduces its GPU frequency  $f_{\text{core}}^{\text{L}}$  by 167 MHz. We find that reaching the limit temperature will reduce the load put on the processor by the detection mechanism and reduce its temperature which lead to a lower fault injection rate.

#### D. Temperature-based guardband analysis

The temperature-based guardband analysis helps identifying voltage margins of the Raspberry Pi models. While this analysis supports the cloud provider in selecting an undervolt offset, its core principle can also be exploited by users to uncover the scrooge attack. This benchmark consists of three stages: 1) booting the operating system while undervolted before 2) adjusting the SoC's temperature either actively or passively and 3) running a billion iterations of the multiplication benchmark described in §IV as a single-threaded process. We set the CPUFreq governor to *performance* right before starting the multiplication benchmark. This will guarantee that the multiplication benchmark is started at a well defined temperature and that it runs at a constant, maximum frequency. Once the benchmark is finished we repeat the following process, in which we reduce the ARM CPU voltage level in the configuration before rebooting, until the Raspberry Pi no longer boots because the supply voltage has gone below threshold.

The results of this analysis are shown in Figure 4. All Raspberry Pi models keep a sufficient margin with their nominal voltage (connected black bullets) configuration to the critical region. Further undervolting of the ARM CPU into the critical region results in occasionally failing processes. Undervolting the ARM CPU beyond threshold voltage makes it impossible to boot the hardware. Our multiplication benchmark, that verifies the correct operation of the CPU, never detected an incorrect result. We explain this characteristic of the multiplication benchmark, which is purely based on arithmetic operations, by not being on a timing-critical path to force an incorrect operation of the CPU.

We can also see that the undervolting depends directly on the SoC's temperature. For instance on the Raspberry Pi 3B we clearly observe slightly rising regions, which result from small adjustments made by the AVS system. This is mainly due to the resistivity of the circuitry that increases with temperature.




Figure 4: Temperature-dependent guardband measurements of latest Raspberry Pi models B. Triangles indicate lower (▽) and upper (△) frontier measurements while bullets (●) indicate nominal measurements.




Figure 5: Temperature-dependent failure rate of Raspberry Pi models at different undervolt levels.

The Raspberry Pi 4B can only be undervolting once to level  $-1$  at  $-15$  mV due to missing overclocking <sup>2</sup> support in the firmware. This undervolt limit is indicated by the blue line. However, we observe some basic overheating protection mechanism that slightly lowers nominal voltage by  $-12.5$  mV in the range of  $50$  °C to  $70$  °C.

#### E. Implications on the detection method success rate

Overall we ran 741 guardband analyses of which 265 had failed while operating in the critical region. Among the 265 failed runs we identified 407 process failures of the following 5 types: (1) NULL pointer dereferences (20.3%), (2) paging requests (46.4%), (3) read from unreadable memory (5.4%), (4) write to read-only memory (0.9%), and (5) freeze during boot (26.7%). These types of failures, with the exception of (5), usually generate a kernel oops. A kernel oops happens when the operating system kernel detects an incorrect behavior of a process and can possibly resume execution of the system. In some cases execution cannot be resumed because of system dependencies or unavailable system resources as a result of the failing process. The kernel will raise a panic and halt the system if a kernel oops occurs in an interrupt handler.

The ARM architecture provides the Exception Syndrome Register [45], which the operating system can consult to diagnose the type of exception generated by a process. If possible, the operating system kernel will log this kernel

oops diagnose in the system log. We used this information to analyze the guardband failures and summarized it in Figure 5. Notice that the 4B is not included, as it’s firmware does not provide undervolting support and we could not provoke any failures in this system. Our analysis indicates that at  $60$  °C we have the highest probability with a 40 % chance on the 3B+ respectively a 90 % chance on the 3B to provoke a failure in a system operating in the critical region. With an even more aggressive undervolting at  $-100$  mV, failures can already be reliably provoked at  $40$  °C on both devices. These failures were provoked in at least 33 different processes (34 % user, 15 % kernel processes and 51 % unknown processes) of which the multiplication benchmark is never among the known failed processes. For the 3B+ the failure rate is dropping at  $70$  °C due to the temperature soft-limit throttling in the firmware to bring the system back into the safe region. At nominal voltage no failures could be provoked in any system.

#### F. Energy efficiency and “reliability”

Despite stressing the systems for several hours, we could not provoke any failures in these systems using STRESS-NG. We show in heat map Table IV the energy efficiency for all three Raspberry Pi models based on the ETR ratio of the undervolting to the nominal setup. For the measurements we use two cooling setups: active and passive cooling. We ran up to 169 stressors sequentially of which 27 are shown in the heat map. Each stressor was configured with a timeout of 60s.

Our results show that when the Raspberry Pi’s are actively cooled, we can achieve higher energy efficiency. It is even possible to undervolt the device further. For example, with the 3B+ we were able to undervolt up to  $-100$  mV and in some rare instances also run the benchmark successfully. Occasionally we observe minimally better energy efficiency in the passive cooling setup than in the active cooling setup (e.g., futex with  $-34$  %,  $-3$  % and  $-16$  % and  $2$  %,  $-1$  % and  $-2$  %). We noticed that some stressors have a large variance in the number of operations. Hence, if these stressors achieve a higher than average number of operations during the measurement, their energy efficiency improves proportionally. With the 4B we notice only minor improvements in energy efficiency. Again, this is due to the lack of overclocking support but also because of the lower core supply voltage compared to the other models. On average across the 27

<sup>2</sup><https://www.raspberrypi.org/documentation/configuration/config-txt/overclocking.md> Last accessed on 2021-04-23

Table IV: STRESS-NG ETR heat map indicating the relative energy efficiency for an undervolted setup compared to a nominal setup. The darker the shade, the more energy-efficient the stressor ran.

	Cooling Model	Undervolt	Stressors																										
			aito	atomic	bsearch	clock	fork	futex	get	hrtimers	hsearch	icache	judy	kcmp	kill	lsearch	membarrier	mergesort	msg	pipe	poll	sem	sigsegv	sysfs	timer	tsearch	urandom	vfm-rw	wcs
active	3B	-75 mV	0.94	0.95	0.96	0.95	0.92	1.02	1.03	0.90	0.95	0.95	0.99	0.93	0.93	0.94	1.02	0.94	0.91	0.96	0.95	0.93	0.94	0.91	0.94	0.99	0.95	0.96	0.94
	3B+	-75 mV	0.89	0.94	0.93	0.93	0.87	0.99	0.94	1.00	0.95	0.94	1.01	0.93	0.96	0.94	0.97	0.94	0.95	0.92	0.95	0.93	0.93	0.92	0.94	0.95	0.95	0.97	0.94
	4B	-15 mV	1.01	0.99	1.02	0.99	1.02	0.98	1.00	1.06	1.00	0.98	0.96	1.00	1.04	0.99	0.96	1.00	0.97	0.70	0.98	0.98	0.99	1.00	0.97	0.91	0.98	1.00	0.99
passive	3B	-75 mV	0.88	0.95	0.92	0.93	0.91	0.66	0.76	0.63	0.94	0.94	0.97	0.93	0.94	0.94	0.93	0.93	0.92	1.03	0.93	0.95	0.92	0.95	0.92	0.96	0.94	0.96	0.93
	3B+	-75 mV	0.95	0.95	0.96	0.95	0.98	0.97	0.99	0.80	0.95	0.95	0.94	0.95	0.95	0.95	0.98	0.95	0.95	0.96	0.94	0.95	0.94	0.91	0.95	0.97	0.96	0.97	0.95
	4B	-15 mV	1.00	0.97	1.02	0.99	1.01	0.84	1.00	1.10	0.99	1.03	0.99	0.98	1.00	1.00	0.97	1.00	1.03	1.01	1.00	1.05	1.04	0.87	1.00	0.91	0.99	0.97	0.99

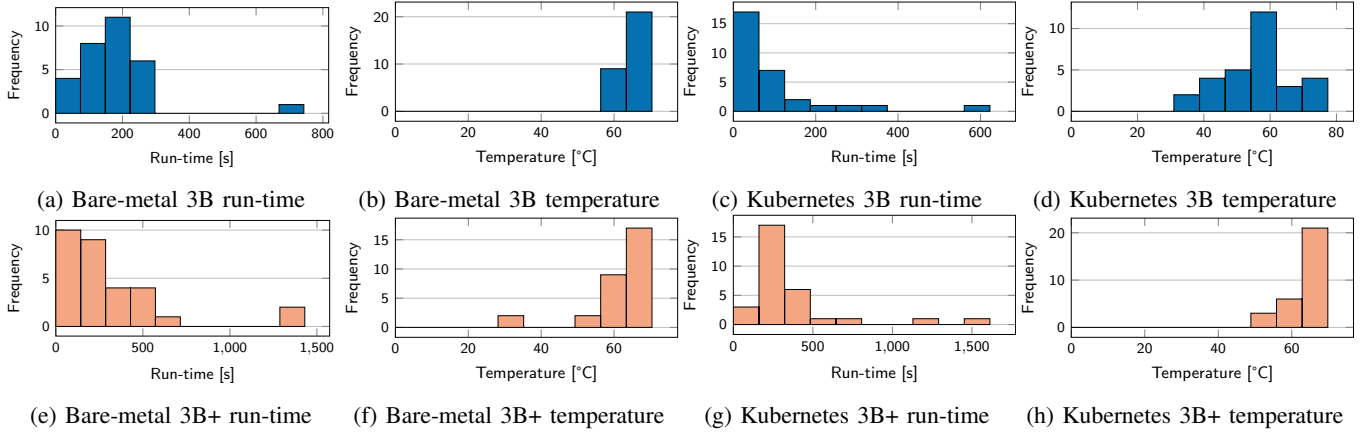


Figure 6: Run-time and temperature histograms of bare-metal and container instances

stressors 5% / 9% (active/passive) were saved on the 3B, 6% / 6% were saved on the 3B+, and 2% / 1% were saved on the 4B. The highest energy efficiency observed on the 3B was -10% / -37% on the hrtimers stressor. On the 3B+ -13% / -20% were saved on the fork / hrtimers stressor. Finally, the -30% / -16% were saved on the 4B with the hrtimers / futex stressor.

### G. Detection method parameters

In this subsection we quantify the detection method parameters (*i.e.*, run-time and temperature) based on undervolted bare-metal and container instances. Deploying virtual machines on the Raspberry Pi is impracticable and were therefore not included in our evaluation. To run containers on the Raspberry Pi we deployed a small Kubernetes cluster.

Figure 6 shows histograms with crashes on bare-metal and container instances deployed on the 3B and 3B+. We show the run-time of our detection method and the temperature at which instances crashed. Our observations made with the temperature-based guardband analysis in subsection V-E are confirmed by the temperature histograms. The run-time strongly depends on the processor’s capability to heat up to a certain temperature and is therefore not an ideal parameter. We observe clear differences between the thermal designs of the two models. For the 3B our detection method requires about 175 s / 30 s to reach 62 °C to crash bare-metal or container instances. On the 3B+ we require about 145 s / 250 s to reach

62 °C to crash bare-metal or container instances. Interestingly, container instances crash on the 3B earlier than bare-metal instances. We assume the computing requirements from the container environment work in favor of the detection method.

## VI. DISCUSSION

From our evaluation we conclude that the detection method is best used in combination with other processes such as in STRESS-NG. The user even has the option to scale the number of threads in the detection method to adjust the crash time of an instance as well as the injection rate. A simple CPU-bound program like the multiplication benchmark turns out to be ideal for injecting faults in an undervolted setup. The advantage of such a simple CPU-bound program is that it is unlikely to inject faults during its own execution and can run until a kernel panic while raising heat dissipation. In terms of energy efficiency we observed that by undervolting the cloud provider can save on average 5% and up to 37% for specific workloads on ARM processors.

**RAI:** as shown by our extensive experimental evaluation, in order to pull off a stealthy undervolting strategy, a malicious cloud provider must exchange any firmware configuration to undervolt the hardware and intercept any voltage requests coming from users.



**RA2:** a cloud user can uncover such an undervolting strategy by running a simple CPU-bound benchmark until enough processes have failed to render the cloud instance unavailable. The drawback of this detection method is that it is non-selective and cloud instances can fail either soon or late.

## VII. RELATED WORK

Undervolting the supply voltage for energy savings has been explored on CPUs for ARM [46, 47], x86 processors [16, 48] the Itanium micro-architecture [49], and for POWER-7 processors [36]. This experimental undervolting approach has been extended to GPUs [50] and FPGAs [40] as well. On the CPU side, frameworks to automate and optimize the process of undervolting have been developed [14, 46]. Recently, AMD has announced an undervolting product/framework for their most recent Ryzen 5000 CPUs [51]. In [52] the authors discuss the trade-off between the reduced energy cost and the SLA violation penalties introduced by higher node failures of undervolting X86 and ARM nodes. In CLKSCREW [20], the undervolting capabilities of modern ARM processors is exploited to compromise system security, by targeting undervolting faults to specific hardware components to extract cryptographic keys.

## VIII. CONCLUSION AND OPEN CHALLENGES

A cloud provider can obfuscate the undervolting of processors and even run workloads up to 37% more energy-efficiently. However, by undervolting its infrastructure, the cloud provider incurs a major risk. Not only does the cloud provider reduce the margin of error but also the system's stability is at stake. Cloud users can with high probability detect such situations and exploit them using a simple CPU-bound benchmark. To some extent, the cloud provider can mitigate stability issues with appropriate cooling systems. However, it is questionable if the gains of undervolting the infrastructure outweigh the costs of such cooling systems.

Cloud users' options to detect an undervolting ARM instance remain limited and, as shown in this paper, essentially depend on the probability to inject faults non-selectively in processes. As our temperature-based guardband analysis and failure evaluation have shown, the higher the processor's temperature, the more likely faults can be injected into processes. Despite such a powerful cloud provider attacker model, cloud users have an exploitable weak link. Their only option for presuming a potentially undervolting instance is by increasing the processor's heat dissipation. Heat dissipation is increased by tuning the CPU frequency and load to the processor's limit. Under these thermal conditions and an undervolting setup the fault injection probability in processes is rising. Ideally cloud instances will become unavailable and violate the SLA as a result of continuously failing processes. Our detection method depends strongly on hardware and how systems such as firmware and AVS react to excessive heat dissipation. As future plans, we intend to expand this study to a more diverse set of ARM-based hardware targets, focusing in particular on current and future cloud offerings. We would also like to make

our detection method more deterministic by injecting faults in processes more selectively.

## ACKNOWLEDGMENTS & DISCLAIMER

The views and opinions of the authors do not necessarily reflect those of the U.S. government or Lawrence Livermore National Security, LLC neither of whom nor any of their employees make any endorsements, express or implied warranties or representations or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of the information contained herein. This work was partially prepared by LLNL under Contract DE-AC52-07NA27344 (LLNL-CONF-817551) and by the European Union's Horizon 2020 research and innovation programme under the LEGaTO Project ([legato-project.eu](http://legato-project.eu)), grant agreement No 780681.

## REFERENCES

- [1] *Ampere eMAG 8180 64-bit Arm Processor*, Amp 2018-0007 ed., Ampere Computing, 4655 Great America Parkway, Suite 601, Santa Clara, CA 95054, 2018.
- [2] "Ampere Altra: The World's First Cloud Native Processor," <https://amperecomputing.com/altra/>, Nov 2020, last accessed on 2021-04-23.
- [3] "Huawei Unveils Industry's Highest-Performance ARM-based CPU," <https://www.huawei.com/en/news/2019/1/huawei-unveils-highest-performance-arm-based-cpu>, Jan 2019, last accessed on 2021-04-23.
- [4] "NVIDIA Grace CPU," <https://www.nvidia.com/en-us/data-center/grace-cpu/>, Apr 2021, last accessed on 2021-04-23.
- [5] "AWS Graviton Processor," <https://aws.amazon.com/ec2/graviton/>, last accessed on 2021-04-23.
- [6] J. Barr, "Coming Soon - Graviton2-Powered General Purpose, Compute-Optimized, & Memory-Optimized EC2 Instances," <https://aws.amazon.com/de/blogs/aws/coming-soon-graviton2-powered-general-purpose-compute-optimized-memo>, Dec 2019, last accessed on 2021-04-23.
- [7] *ThunderX Family of Workload Optimized Processors*, Cavium, 2315 N. First Street, San Jose, CA 95131, 2016.
- [8] T. Yoshida, "Fujitsu High Performance CPU for the Post-K Computer," [https://old.hotchips.org/hc30/2conf/2.13\\_Fujitsu\\_HC30.Fujitsu.Yoshida.rev1.2.pdf](https://old.hotchips.org/hc30/2conf/2.13_Fujitsu_HC30.Fujitsu.Yoshida.rev1.2.pdf), Aug 2018, last accessed on 2021-04-23.
- [9] "Raspberry Pi Products," <https://www.raspberrypi.org/products/>, last accessed on 2021-04-23.
- [10] Y. Léger, "Public Preview," <https://blog.scaleway.com/online-labs-public-preview>, Oct. 2014, last accessed on 2021-04-23.
- [11] "Neoverse N1," <https://developer.arm.com/ip-products/processors/neoverse/neoverse-n1>, last accessed on 2021-04-23.
- [12] T. Burd and R. Brodersen, "Energy efficient cmos microprocessor design," in *2014 47th Hawaii International Conference on System Sciences*, vol. 1. Los Alamitos, CA, USA: IEEE Computer Society, Jan 1995, p. 288. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/HICSS.1995.375385>
- [13] V. M. van Santen, H. Amrouch, N. Parihar, S. Mahapatra, and J. Henkel, "Aging-aware voltage scaling," in *2016 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2016, pp. 576–581.
- [14] K. Parasyris, P. Koutsovasilis, V. Vassiliadis, C. D. Antonopoulos, N. Bellas, and S. Lalis, "A framework for evaluating software on reduced margins hardware," in *2018 48th Annual*

- IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2018, pp. 330–337.
- [15] G. Papadimitriou, A. Chatzidimitriou, D. Gizopoulos, V. J. Reddi, J. Leng, B. Salami, O. S. Unsal, and A. C. Kestelman, “Exceeding conservative limits: A consolidated analysis on modern hardware margins,” *IEEE Transactions on Device and Materials Reliability*, vol. 20, no. 2, pp. 341–350, 2020.
- [16] P. Koutsovasilis, K. Parasyris, C. D. Antonopoulos, N. Bellas, and S. Lalis, “Dynamic undervolting to improve energy efficiency on multicore x86 cpus,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 12, pp. 2851–2864, 2020.
- [17] K. Murdock, D. Oswald, F. D. Garcia, J. Van Bulck, D. Gruss, and F. Piessens, “Plundervolt: Software-based Fault Injection Attacks against Intel SGX,” in *Proceedings of the 41st IEEE Symposium on Security and Privacy (S&P’20)*, 2020, 41st IEEE Symposium on Security and Privacy (S&P’20).
- [18] Z. Kenjar, T. Frassetto, D. Gens, M. Franz, and A.-R. Sadeghi, “V0ltpwn: Attacking x86 processor integrity from software,” in *29th {USENIX} Security Symposium ({USENIX} Security 20)*, 2020, pp. 1445–1461.
- [19] P. Patel, A. H. Ranabahu, and A. P. Sheth, “Service level agreement in cloud computing,” 2009.
- [20] A. Tang, S. Sethumadhavan, and S. Stolfo, “{CLKSCREW}: exposing the perils of security-oblivious energy management,” in *26th {USENIX} Security Symposium ({USENIX} Security 17)*, 2017, pp. 1057–1074.
- [21] Z. Chen, G. Vasilakis, K. Murdock, E. Dean, D. Oswald, and F. D. Garcia, “VoltPillager: Hardware-based fault injection attacks against intel SGX Enclaves using the SVID voltage scaling interface.” USENIX Association, Aug. 2021. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity21/presentation/chen-zitai>
- [22] L. Upton, “Raspberry Pi colocation,” <https://www.raspberrypi.org/blog/raspberry-pi-colocation>, Apr 2013, last accessed on 2021-04-23.
- [23] E. Blem, J. Menon, and K. Sankaralingam, “Power struggles: Revisiting the RISC vs. CISC debate on contemporary ARM and x86 architectures,” in *2013 IEEE 19th International Symposium on High Performance Computer Architecture (HPCA)*, 2013, pp. 1–12.
- [24] H. Esmaeilzadeh, E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger, “Dark silicon and the end of multicore scaling,” in *2011 38th Annual International Symposium on Computer Architecture (ISCA)*, 2011, pp. 365–376.
- [25] E. Le Sueur and G. Heiser, “Dynamic Voltage and Frequency Scaling: The Laws of Diminishing Returns,” in *Proceedings of the 2010 International Conference on Power Aware Computing and Systems*, ser. HotPower’10. USA: USENIX Association, 2010, p. 1–8.
- [26] M. Eleršič, “linux-intel-undervolt,” <https://github.com/mihic/linux-intel-undervolt>, Aug 2017, last accessed on 2021-04-23.
- [27] D. Brodowski, “Linux CPUFreq - CPU frequency and voltage scaling code in the Linux kernel,” <https://www.kernel.org/doc/html/latest/cpu-freq/index.html>, last accessed on 2021-04-23.
- [28] M. Weiser, B. Welch, A. Demers, and S. Scott, “Scheduling for Reduced CPU Energy,” in *First Symposium on Operating Systems Design and Implementation (OSDI 94)*. Monterey, CA: USENIX Association, Nov. 1994. [Online]. Available: <https://www.usenix.org/conference/osdi-94/scheduling-reduced-cpu-energy>
- [29] J.-T. Wamhoff, S. Diestelhorst, C. Fetzer, P. Marlier, P. Felber, and D. Dice, “The TURBO Diaries: Application-Controlled Frequency Scaling Explained,” in *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, ser. USENIX ATC’14. USA: USENIX Association, 2014, p. 193–204.
- [30] “Frequently Asked Questions about Enhanced Intel SpeedStep Technology for Intel,” <https://www.intel.com/content/www/us/en/support/articles/000007073/processors.html>, Jun. 2020, last accessed on 2021-04-23.
- [31] *Cool’n’Quiet Technology Installation Guide for AMD Athlon 64 Processor Based Systems*, 0th ed., Advanced Micro Devices Inc., Jun. 2004.
- [32] *AMD PowerNow! Technology*, A ed., Advanced Micro Devices Inc., Nov. 2000.
- [33] L. S. Nielsen, C. Niessen, J. Sparso, and K. Van Berkel, “Low-power operation using self-timed circuits and adaptive scaling of the supply voltage,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 2, no. 4, pp. 391–397, 1994.
- [34] J. L. Henning, “SPEC CPU2006 benchmark descriptions,” *ACM SIGARCH Computer Architecture News*, vol. 34, no. 4, pp. 1–17, 2006.
- [35] C. Bienia, S. Kumar, J. P. Singh, and K. Li, “The PARSEC Benchmark Suite: Characterization and Architectural Implications,” in *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques*, October 2008.
- [36] Y. Zu, C. R. Lefurgy, J. Leng, M. Halpern, M. S. Floyd, and V. J. Reddi, “Adaptive guardband scheduling to improve system-level efficiency of the POWER7+,” in *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2015, pp. 308–321.
- [37] L. Tan, N. DeBardeleben, Q. Guan, S. Blanchard, and M. Lang, “Using virtualization to quantify power conservation via near-threshold voltage reduction for inherently resilient applications,” *Parallel Computing*, vol. 73, pp. 3–15, 2018, parallel Programming for Resilience and Energy Efficiency. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167819117300996>
- [38] Y. Kim, L. K. John, S. Pant, S. Manne, M. Schulte, W. L. Bircher, and M. S. S. Govindan, “Audit: Stress Testing the Automatic Way,” in *2012 45th Annual IEEE/ACM International Symposium on Microarchitecture*, 2012, pp. 212–223.
- [39] Z. Hadjilambrou, S. Das, M. A. Antoniadis, and Y. Sazeides, “Sensing CPU Voltage Noise Through Electromagnetic Emanations,” *IEEE Computer Architecture Letters*, vol. 17, no. 1, pp. 68–71, 2018.
- [40] B. Salami, E. B. Onural, I. E. Yuksel, F. Koc, O. Ergin, A. C. Kestelman, O. S. Unsal, H. Sarbazi-Azad, and O. Mutlu, “An Experimental Study of Reduced-Voltage Operation in Modern FPGAs for Neural Network Acceleration,” in *2020 50th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2020.
- [41] K. Beer, “How encryption works in AWS,” <https://image.slidesharecdn.com/repeat-how-encryption-works->, Jun. 2019, last accessed on 2021-04-23.
- [42] “Open Portable Trusted Execution Environment,” <https://www.op-tee.org>, last accessed on 2021-04-23.
- [43] “AWS Nitro System,” <https://aws.amazon.com/ec2/nitro/>, last accessed on 2021-04-23.
- [44] *PowerSpy2*, 1st ed., Alciom, 4 Mar. 2013.
- [45] *Arm Architecture Reference Manual: Armv8, for Armv8-A architecture profile*, Arm ddi 0487f.b (id040120) ed., Arm Limited, Mar. 2020.
- [46] G. Papadimitriou, M. Kaliorakis, A. Chatzidimitriou, D. Gizopoulos, P. Lawthers, and S. Das, “Harnessing voltage margins for energy efficiency in multicore cpus,” in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*, 2017, pp. 503–516.
- [47] G. Papadimitriou, A. Chatzidimitriou, M. Kaliorakis, Y. Vastakis, and D. Gizopoulos, “Micro-viruses for fast system-level voltage margins characterization in multicore cpus,” in *2018 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 2018, pp. 54–63.

- [48] G. Papadimitriou, M. Kaliorakis, A. Chatzidimitriou, C. Magdalinos, and D. Gizopoulos, "Voltage margins identification on commercial x86-64 multicore microprocessors," in *2017 IEEE 23rd International Symposium on On-Line Testing and Robust System Design (IOLTS)*. IEEE, 2017, pp. 51–56.
- [49] A. Bacha and R. Teodorescu, "Dynamic reduction of voltage margins by leveraging on-chip ecc in itanium ii processors," in *Proceedings of the 40th Annual International Symposium on Computer Architecture*, 2013, pp. 297–307.
- [50] J. Leng, A. Buyuktosunoglu, R. Bertran, P. Bose, and V. J. Reddi, "Safe limits on voltage reduction efficiency in gpus: a direct measurement approach," in *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2015, pp. 294–307.
- [51] I. Cutress, "AMD Precision Boost Overdrive 2: Adaptive Undervolting For Ryzen 5000 Coming Soon," <https://www.anandtech.com/show/16267/amd-precision-boost-overdrive-2-adaptive-undervolting-for-ryzen-5000-coming-soon>, Nov 2020, last accessed on 2021-04-23.
- [52] C. Kalogirou, P. Koutsovasilis, C. D. Antonopoulos, N. Bellas, S. Lalis, S. Venugopal, and C. Pinto, "Exploiting cpu voltage margins to increase the profit of cloud infrastructure providers," in *2019 19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*, 2019, pp. 302–311.