

# The Forgotten Threat of Voltage Glitching: A Case Study on Nvidia Tegra X2 SoCs

Otto Bittner<sup>\*§</sup>, Thilo Krachenfels<sup>\*§</sup>, Andreas Galauner<sup>†</sup> and Jean-Pierre Seifert<sup>\*‡</sup>

<sup>\*</sup> Technische Universität Berlin, Chair of Security in Telecommunications

<sup>†</sup> Independent Researcher

<sup>‡</sup> Fraunhofer SIT

**Abstract**—Voltage fault injection (FI) is a well-known attack technique that can be used to force faulty behavior in processors during their operation. *Glitching* the supply voltage can cause data value corruption, skip security checks, or enable protected code paths. At the same time, modern systems on a chip (SoCs) are used in security-critical applications, such as self-driving cars and autonomous machines. Since these embedded devices are often physically accessible by attackers, vendors must consider device tampering in their threat models. However, while the threat of voltage FI is known since the early 2000s, it seems as if vendors still *forget* to integrate countermeasures. This work shows how the entire boot security of an *Nvidia SoC*, used in Tesla’s autopilot and Mercedes-Benz’s infotainment system, can be circumvented using voltage FI. We uncover a hidden bootloader that is only available to the manufacturer for testing purposes and disabled by fuses in shipped products. We demonstrate how to re-enable this bootloader using FI to gain code execution with the highest privileges, enabling us to extract the bootloader’s firmware and decryption keys used in later boot stages. Using a hardware implant, an adversary might misuse the hidden bootloader to bypass trusted code execution even during the system’s regular operation.

**Index Terms**—fault injection, voltage glitching, SoC, root of trust

## I. INTRODUCTION

Modern systems on a chip (SoCs) are used in various applications where chip security directly relates to human safety, such as self-driving cars and autonomous machines. To provide secure and authenticated operation of the system, SoCs contain a trust anchor in the form of a secured and tamper-proof bootloader. That bootloader’s task is to allow the execution of only authenticated firmware components, protect intellectual property (IP), set security configurations, and transfer control to less privileged boot stages. Typically, dedicated processor cores on the SoC act as the root of trust (RoT) and are responsible for these delicate tasks.

In order to function as intended, integrated circuits (ICs) need to be operated under specified conditions, for instance, within the rated supply voltage, clock stability, temperature, and electromagnetic field ranges [1]. This dependency can be misused to force faulty behavior during the chip’s operation. Hence, the susceptibility of electronic circuits to coincidentally or deliberately injected faults has been studied for some decades. Especially in the smartcard field, fault injection (FI) attacks used to extract secrets from cryptographic algorithms

were investigated and counteracted around the year 2000 [1]–[3]. Next to the intended corruption of data values, faults can be used to skip security checks, enter protected code paths, or gain code execution [4], [5]. During the past years, attacks against microcontrollers and SoCs using laser-based [6] and electromagnetic [7], [8] FI have been presented. While these techniques offer high accuracy in targeting a specific part of the chip, they also require comparatively sophisticated setups.

A simpler approach to inject faults into the system is voltage FI, where the supply voltage is over- or undervolted for a short moment [9]. The technique is also referred to as *voltage glitching*. Even though the technique is simpler to execute, it is repeatedly used to attack modern targets, like the Nvidia Tegra X1 SoC [10]. By injecting faults into security registers, the bootloader’s code was extracted and used to find a firmware bug enabling unauthenticated code execution on the system. The fact that such simple attacks are still possible, 20 years after smartcards have been hardened against FI, suggests that chip manufacturers seem to have ignored this threat. Although protecting SoCs can be more difficult due to multiple power domains, complex power trees, and higher power consumption, manufacturers should implement protections against obvious and known attacks that can break the security of the system or even the entire device family.

The subsequent generation of the above mentioned SoC, the Nvidia Tegra X2 (codename “Parker”), is used in safety-critical applications, such as for the Nvidia autonomous driving units DRIVE PX Parker AutoChauffeur and AutoCruise [11] used in Tesla cars [12], or the infotainment system in Mercedes [13] and Hyundai [14] cars. In our work, we address the following question: *Is the Tegra X2 SoC susceptible to voltage FI as well, or was it a bad coincidence that the previous generation was vulnerable?*

**Our contribution** In this work, we indeed show that the processor acting as RoT on the Tegra X2 SoC is susceptible to voltage FI. We demonstrate how an attacker can gain code execution in the secure zone of the boot processor with only cheap and readily available equipment. This capability allows us to extract the content of the internal read-only memory (iROM), containing the first bootloader and key material used for decrypting the code of later boot stages. This endangers the IP of original equipment manufacturers (OEMs) and can defeat trusted code execution. To this end, we explain how a hardware implant can permanently manipulate the RoT. Since

<sup>§</sup>These authors contributed equally to this work.

our attack cannot be easily prevented by firmware patches, we propose and discuss potential mitigations against voltage FI attacks for future chip generations.

**Responsible Disclosure** We responsibly disclosed our findings to Nvidia, including our experimental setup and parameters. Nvidia reconstructed our experiments and confirmed that fault injection impacts the tested Tegra Parker SoC and earlier chips. According to them, all newer Tegra SoCs would contain countermeasures to mitigate these types of attacks. Furthermore, they proposed countermeasures to reduce the effectiveness of voltage fault injection on vulnerable chips, which we discuss in Section VII-B1.

## II. BACKGROUND – VOLTAGE FAULT INJECTION

ICs need to be operated under the specified conditions to function as intended, e.g., within rated supply voltage, clock stability, temperature, and electromagnetic field ranges [1]. This dependency can be misused to force faulty behavior during the chip’s operation. Short supply voltage variations, introduced by glitches on the supply voltage line, can produce computational errors in CMOS circuits [9]. Examples of such errors are memory bit flips, corrupted instructions, and jumping over instructions in a microprocessor. If these errors are forced during the execution of cryptographic algorithms, information about the secret key or plaintext might be leaked [1]–[3]. On the other hand, faults can be used to skip security checks, enter protected code paths, or gain code execution [4], [5]. Voltage FI is a well-studied field, especially due to the low cost of setups. Open-source frameworks, such as the ChipWhisperer [15], provide easy access to both hardware and software to conduct attacks. A recent study shows that the shape of the voltage glitch can improve the attack performance, i.e., reduce the time until a successful glitch is observed [16].

Depending on the design of the target, different approaches can be used to inject faults into the supply voltage rail. If the voltage is supplied externally to the printed circuit board (PCB), an external power supply can introduce glitches through that interface. If the voltage is generated directly on the PCB using a voltage regulator, the injection of glitches becomes more complex. On the other hand, on more advanced systems, such as SoCs, the voltage regulators typically offer communication interfaces to adjust the voltage on demand. These interfaces, if not adequately protected, can be leveraged to inject voltage glitches [17]. In some cases, the interface is even accessible via software [18], [19].

Another alternative is to inject glitches using a so-called crowbar circuit. The idea is to create a short circuit between the voltage line and GND, effectively enforcing a voltage drop [20]. Fig. 1 shows a schematic of such a setup. A transistor acting as a switch – typically an n-channel metal–oxide–semiconductor field-effect transistor (MOSFET) – is connected between the supply voltage input (VCC) of the device under test (DUT) and GND. To reduce noise on the supply voltage rail, PCB designers place so-called decoupling capacitors close to the DUT. Their connection to VCC offers

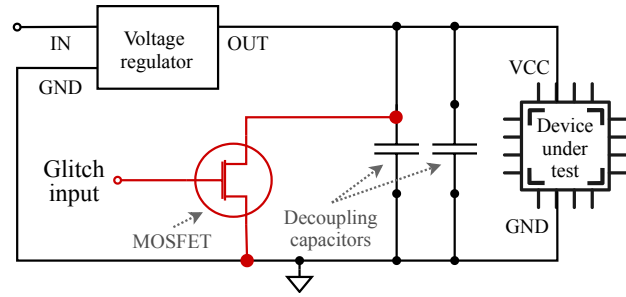


Figure 1: Schematic of a crowbar circuit (red) used to inject supply voltage glitches.

a good point for soldering the MOSFET. As the decoupling capacitors might reduce the effectiveness of the voltage drop, desoldering them can be beneficial to achieve shorter glitches. During the glitch, a high short circuit current will be flowing through the MOSFET, effectively pulling the VCC voltage close to the GND level. It should be noted that modern SoCs typically have more than one power domain, which can complicate finding the correct rail to inject the glitch.

## III. RELATED WORK

As mentioned in the introduction, there is a multitude of examples for successful attacks on computational systems using fault injection. In the following, we will introduce some of those examples that relate to our work more closely.

### A. Impactful Fault Injection

With the continuous digitalization of cars, a huge market that requires secure embedded devices is the automotive sector. In 2018, researchers from Riscure and VU University Amsterdam showed how they successfully extracted firmware from secured car electronic control units (ECUs) using voltage glitching [21]. Furthermore, the study shows how the recovered firmware blobs can be emulated to quickly gain insights into the inner working of the firmware and potentially find vulnerabilities that can be exploited from software. Similarly, the authors of [22] investigated vulnerabilities in microcontroller bootloaders that can be exploited by voltage glitching. Using dynamic and static analysis of the targeted binary, they show how even a multi-glitch attack can be applied on a common microcontroller’s bootloader.

In the past, also gaming consoles have been subject to repeated scrutiny from hacker groups. There exist exploits to gain code execution on both the PlayStation 3 [23] and the Xbox 360 [24]. Remarkably enough, for both consoles, the respective hack leverages fault injection in some way. For the PS3, a write to a memory bus was manipulated to skip the de-allocation of a particular memory region. This could subsequently be leveraged to gain read/write access to the hypervisor’s page table, which gives the attacker full control over the system. In the case of the Xbox 360, the attackers realized that the device’s processor does not fully reset if the reset pin is only asserted for a very short amount of time. This insight was used to skip the signature check for the second

bootloader stage by pulsing the reset pin for 100 ns. Both of these examples showcase how otherwise very tight security measures can be bypassed entirely using fault injection.

### B. Previous Tegra Generations

When the gaming company Nintendo released their newest handheld gaming console called “Nintendo Switch” in 2017, initial teardowns and analyses of the hardware suggested that the SoC employed on the platform could be an Nvidia Tegra X1 [25]. As with previous gaming consoles, the Switch was investigated by several hacker groups. Their research aimed to gain the capability to execute code on the platform, and therefore, to bypass the mandatory code signing required by Nintendo. While the researchers exploited typical software targets like the embedded Webkit-based browser and interfaces to the operating system kernel at first, the BootRom (BR), which enforces the root of trust, later became a target as well [26].

Compromising this part of the boot chain gives an attacker full control over the following boot process, rendering any cryptographic security measures ineffective. Additionally, the BR is not patchable after production, making it impossible for Nintendo or Nvidia to fix discovered flaws. Since the BR is IP of Nvidia, it is not publicly available. Therefore, the BR first needed to be leaked from the SoC for further analysis. By default, due to the read protection, this is not possible. Just before execution is passed from the BR to the first cryptographically verified bootloader, the BR activates a read protection for itself. Multiple independent security researchers successfully applied voltage glitching to skip the instruction that sets the read-lock bit in a control register [10], [27]. Others used the same technique to skip the signature verification of executed bootloaders entirely [28].

## IV. ATTACK APPROACH

### A. Threat Model

The attacks described in the previous section already show that voltage FI can break consumer products and enable unintended usage. However, these kinds of attacks would even pose a physical threat if conducted on a platform used in security-critical applications. The attacks have in common that the capability to execute unauthenticated code is gained. In the case of an autonomous driving unit, an adversary could alter the firmware to tamper with, e.g., how the car reacts to human obstacles. Even if only the cockpit display is tampered with, wrong speed values could be displayed, potentially endangering passengers and pedestrians.

Building on this threat model, we assume an attacker with physical access who wants to gain capabilities to execute unauthenticated code by applying voltage glitching. The attacker might have several options to achieve this goal. Firstly, they could extract the device’s firmware to search it for software bugs that allow code execution. Secondly, the attacker could fault routines that check the authenticity of code before execution. Finally, they could directly gain code execution with elevated privileges by re-enabling a debug or testing

interface. Although the last two approaches do not allow permanent unauthenticated code execution, the attacker could use a hardware implant to provoke the fault whenever needed. In the following, we describe how an attacker might proceed when trying to conduct either of these three approaches.

### B. Attack Procedure

We identified five steps that likely have to be taken when applying any sort of FI attacks.

**Step 1: Determining the feasibility of FI** Ideally, the attacker has some means of executing code on the target CPU. This can either be achieved by a software exploit that lets the attacker execute arbitrary code in an otherwise locked-down environment or by getting access to development hardware that allows code execution for development purposes. The executed code can then be used to test if glitches applied to multiple different voltage rails, clock inputs, reset lines, or similar external control lines might affect the correctness of instructions executed by the CPU. The easiest way to test for this is to build tight endless-loops that constantly add numbers and output them in short succession. If the code behaves correctly, the arithmetic results printed by the CPU need to be correct all the time. If an applied glitch changes the result of an arithmetic operation, it is a strong indicator that the CPU has experienced a fault caused by the applied glitch.

For being able to inject faults into the DUT, some preparation needs to be done. Depending on the target device, that involves finding the right voltage rail, clock input, or reset pin to apply a glitch. If a voltage rail should be glitched, removing capacitors on that rail to achieve a sharper glitch pulse and attaching devices like a MOSFET for shorting supply rails to GND is necessary, cf. Section II. Using the prepared hardware and software setup, fault injection can now be performed to evaluate its effect on the test code.

### Step 2: Identifying the FI target and a success indicator

After determining the feasibility of FI itself, locations in the actual code that should be glitched need to be identified. Typically, a glitch causes an arithmetic result to be wrong, writebacks to memory or registers to fail, or instructions to be skipped entirely. Depending on how the code was written and how the compiler applied optimizations, routines like cryptographic signature checks can be circumvented by skipping the branch to the code path handling a failed signature verification. If this branch instruction can be skipped, the execution continues in the success path and code execution can be gained although the verification failed. For example, in Listing 1, a glitch could potentially skip the `cbz`-instruction in line 4 to make the CPU jump into the code that failed authentication. If FI should prevent the activation of a memory read-protection, either the instruction setting the protection or the writeback into the control register itself can be faulted.

For detecting if a glitch was successful in an automated glitching setup, an externally available indication is needed. This can be, for instance, a signal available at an external pin or specific content in a log file. In case there is no

```

1  push    {fp, lr}
2  bl     load_further_code
3  bl     sig_verify
4  cbz   r0, sigcheck_failed
5  bl     call_authenticated_code
6  sigcheck_failed:
7  bl     signify_auth_error
8  .hang:
9  b     .hang

```

Listing 1: Pseudocode for a signature check that can potentially be skipped. Note that the call in line 5 never returns.

easily available success criterion, side-channel information, e.g., from the power consumption or electromagnetic (EM) emission, can be used for this purpose.

**Step 3: Finding a trigger signal** Since a glitch needs to target the execution of a particular instruction, the attacker needs some timing anchor to measure the time until the glitch needs to be triggered. This timing anchor can be generated, on the one hand, on external interfaces of the device by the software under attack. For instance, some code paths might need input like USB or Universal Asynchronous Receiver Transmitter (UART) traffic, which needs to be generated by the attacker as well. In these cases, generated inputs can be used as a timing anchor. On the other hand, signals used for the initialization of external memories, modems, or similar external ICs are suitable as an anchor as well. Care must be taken to avoid timing anchors where the time between the anchor and the targeted instruction jitters. For example, a modern SoC contains multiple phase-locked loops (PLLs) that are usually configured as soon as code execution starts. PLLs take a varying time to lock and produce a stable frequency output before execution can switch over to them as the new clock source. In this case, an event happening after the code causing jitter is suited better as a stable timing anchor.

When a glitching attempt fails, the target might behave erratically, crash, or hang without any kind of output. In these cases, or for triggering a new glitching attempt, the target needs to be reset by the glitching hardware.

**Step 4: Finding glitch parameters** After preparing the hardware and acquiring a stable timing reference, the right spot for the glitch needs to be found. This can be achieved by sweeping over a timespan while constantly resetting the target and applying glitches at different times until the glitch can be executed more or less reliably. Depending on the target, the timing reference, and the code path, it may be impossible to achieve a 100% success rate. Success rate in this context means the number of glitches that produce the desired result in relation to the number of tries. However, a success rate of 100% is usually not necessary. Often a low enough time-to-success is sufficient, for which the number of tries per timeframe plays a big role. A glitch that only works in 1% of the cases but can be tried hundreds of times per second, would be favorable to one that works in 10% of the cases but

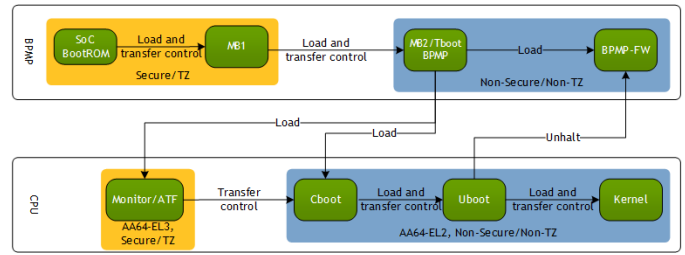


Figure 2: Tegra X2 boot flow. Figure taken from [31].

can only be tried three times per second.

**Step 5: Generating target payload.** When the glitching parameters are determined, the attacker can manipulate or skip an instruction of their will. Depending on their goal, the attacker might have to build custom payloads for the target. For instance, if the attack enables the execution of attacker-controlled unauthenticated code, this code must be created and loaded into the DUT. The payload can carry out different tasks like dumping protected code or secret key material, or chain-loading further code to run more complex payloads.

## V. EXPERIMENTAL SETUP

### A. Device Under Test: Nvidia Tegra X2

The Tegra X2 (model number T186, codenamed "Parker") was introduced in January 2016 to replace the Tegra X1. It is designed for embedded devices requiring high amounts of processing power to support gaming or machine learning applications. It houses a Denver 2 CPU, an ARM Cortex-A57 MPCore CPU, an Nvidia GP10B Pascal GPU, and multiple Cortex-R5 processors for different functionalities. For consumers, the X2 can be acquired from Nvidia as part of the Jetson TX2 module [29]. The TX2 module includes the X2 SoC, as well as external volatile and non-volatile memories, power management controllers, and a Bluetooth and WIFI modem. We used a Jetson TX2 module together with a Jetson TX2 developer kit carrier board [30] for our experiments. The carrier board exposes an array of different signals from the TX2, including the different UART interfaces. To operate the TX2 module, it is plugged on the carrier board, as shown in Fig. 4a, using a 400-pin SAMTEC REF-186137-03 connector.

On the X2 SoC, one of the Cortex-R5 processors, the Boot and Power Management Processor (BPMP), is of particular interest to us as it is responsible for the initial boot process. The BPMP is the first processor in the boot flow and starts execution in the Secure/TZ mode (see Fig. 2). In this mode, the BootRom (BR) and MicroBoot1/nvboot (MB1) are executed. The BR is hardcoded into the X2's iROM and forms the RoT for the SoC. It can only be changed slightly through a patching mechanism that uses fuse storage to apply patches to the BR, called iPatch. However, there is no information publicly available on this feature.

For the X2, the privileges associated with being in Secure/TZ mode are not publicly described by Nvidia either. While the Tegra X1 TRM gives information on how Secure/TZ mode is implemented, the X2 Technical Reference Manual

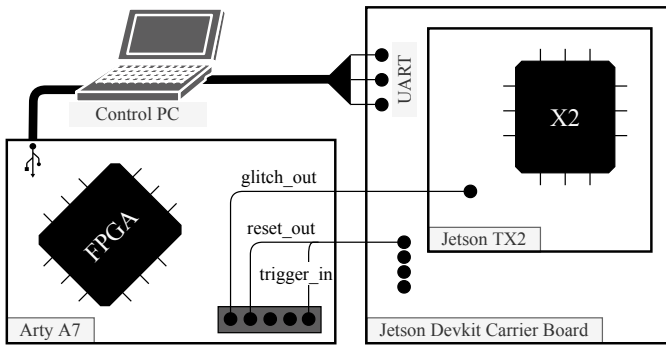


Figure 3: Overview of the FI setup.

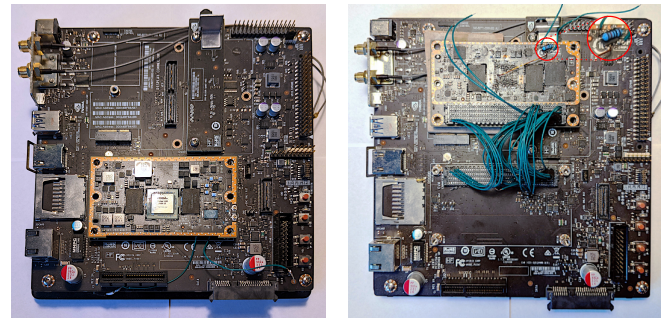
(TRM) [32] does not. From the processor type, we can infer that the Cortex-R5 does not implement ARM’s TrustZone. This aligns with the fact that the BPMP on the X1, an ARM7TDMI processor, did not implement ARM’s TrustZone either. From the previous work on the X1 [10], we can assume that the iROM should only be fully readable in the Secure/TZ mode. However, we only developed a closer understanding of the mode while working with the X2. Our findings are described in the Sections VI-B1 and VI-C2.

After startup, the BR loads the second boot stage called MB1, which is provided by Nvidia as an encrypted and signed binary. It can only be changed by Nvidia, which is advertised as a feature to allow for post-production modifications of the boot flow. MB1 resides in external eMMC memory. Once the BR successfully verifies and decrypts MB1, it yields control to it. During the execution of MB1, more devices and cores are brought up. For OEMs, Nvidia also offers possibilities to implement secure boot features by attesting the integrity of the next boot stage, TegraBoot/TBoot-BPMP (MB2), and possibly encrypting it.

As control is yielded to the MB2 stage, privileges are dropped to the Non-Secure/Non-TZ mode. One task of MB2 is to hand over control to the CPU complex (CCPLEX) (see the lower half in Fig. 2). The boot stages that are executed on the CCPLEX are not relevant for us as they do not have access to the protected iROM of the BPMP. Since MB2 is modifiable by OEMs, no security features are keeping us from changing the MB2 binary and executing our own code. However, as only OEMs are intended to introduce modifications at this level, there is no official documentation on how to build a working binary.

### B. Fault Injection Setup

The three main components of our setup are (a) a control PC, (b) a field-programmable gate array (FPGA), and (c) the X2 SoC, see Fig. 3. The FPGA is a Xilinx Artix-7 XC7A35T, located on a Digilent Arty-A7 development board that provides access to the FPGA’s general-purpose I/O (GPIO) ports and offers a USB-to-UART bridge. Through the UART bridge, we can establish a communication channel between the control PC and the FPGA. The FPGA is used to have exact control



(a) Module without adapter

(b) Module with adapter

Figure 4: TX2 module mounted on the developer kit carrier board.

over the timing of the glitches. Since that timing is on the nanosecond scale, speed is a crucial consideration.

To configure the glitch parameters, we implemented a command-based protocol on the FPGA. This implementation is heavily based on an open-source implementation by chip.fail [33]. The project includes Verilog code for an FPGA and the necessary Python scripts to control the FPGA using the command-based protocol. In our case, the Python code implements a brute force algorithm to randomly traverse the search space of glitch parameters and configure the FPGA for each parameter combination. The glitch parameters are (a) the glitch length and (b) the delay between the trigger signal and the beginning of the glitch activation (offset).

The control PC can also communicate with the X2 DUT using a USB-to-UART adapter, which allows us to read early boot logs and output from programs we run on the X2. For resetting the X2, we use one of the reset signals exposed on the carrier board. By connecting an FPGA GPIO port to the carrier board, we can reset the X2 by toggling the pin.

To inject voltage glitches, we soldered an n-channel Infineon IRF8736PbF MOSFET to the targeted voltage rail on the TX2 module with a pull-down resistor on its gate, see zoom-in of Fig. 4b. The transistor’s gate is controlled by a Maxim Integrated MAX4619 multiplexer that acts as a level shifter and switches between ground and 5 V. The multiplexer, in turn, is controlled by an FPGA GPIO to trigger a glitch.

## VI. RESULTS

### A. Injecting Faults into Code Execution

For testing the BPMP’s susceptibility to FI, we ideally need the ability to execute our own code on this processor, as described by Step 1 in Section IV-B. In our code, we need to issue a stable trigger signal available at an external pin of the DUT. Since we might regularly have to reset the chip during testing, our code should be executed as early as possible in the boot flow. Furthermore, we need to find the correct voltage rail on the PCB that correlates with the BPMP.

1) *MB2 Code Execution*: In Section V-A, we mentioned that the MB2 boot stage is intended to be modified by OEMs. However, no public documentation is available on how to do

this. Thus, we had to develop a working binary for the Cortex-R5 processor blindly.

In order to cross-compile binaries for the BPMP, we used the GNU Arm Embedded Toolchain, version 9.3.1. To set up a C runtime environment for MB2- the `.bss`-section must be zeroized so that uninitialized variables in the C code have a predictable default value. Furthermore, the stack pointer must be set to the correct address before jumping into the main function. This startup code is written in assembly and built using the ARM toolchain mentioned above. We configured the compiler to skip all standard system startup files or libc function imports.

Since we had to define the base address for the binary explicitly, we used a custom linker script. We found the base address by running a tool called `rbasefind` [34] on the original binary. It searches the binary for strings and pointers and calculates the number of pointers directing to the found strings for different base addresses. Depending on the number of pointers that point to a string, the tool recommends the most likely base address candidates. In our case, the tool showed the address `0x52000000` as the most promising candidate. Loading the original binary into Ghidra using this address indicated its correctness. This was confirmed by successful code execution in MB2, which gave us control over the BPMP in the Non-Secure/Non-TZ mode. Most importantly, at this point, we are able to communicate with the control PC via UART and define trigger signals using GPIOs.

2) *Finding the BPMP's Voltage Rail*: Finding the correct voltage rail to inject glitches was the next challenge. There is only a high-level description of the TX2 module available, mentioning information necessary for hardware designers to integrate the module on their own carrier boards. The full power tree is not publicly documented. However, several switching voltage regulators can be identified by their large inductors and capacitors when looking at the module. Three voltage regulators are uP1666Q buck controllers that support the Nvidia OpenVReg Type 2+1 PWMVID feature [35]. They allow dynamic voltage control via a pulse-width modulation (PWM) signal, which the X2 SoC generates.

Since MB1 starts most of the processors in the X2, there needs to be a way of controlling the PWM signals. The way developers can influence the behavior of MB1 is the Boot Configuration Table (BCT). This table includes configuration values for devices that MB1 initializes and is built by a software toolchain included in Nvidia's Linux for Tegra (L4T) package [36]. The Driver Package Development Guide [37] describes how the voltage configuration through the BCT is done in detail. From the L4T package files that control the configuration, we can learn the default voltage levels and within which boundaries the rail should be operated. Since the three rails are configured with different voltage levels, we can map the rail names `VDD_SYS_SOC` (0.95 V), `VDD_SYS_CPU` (0.78 V), and `VDD_SYS_GPU` (0.93 V) to the respective voltage regulators on the PCB. Now knowing which voltage level correlates with which power domain, we can probe the decoupling capacitors present on the backside of the PCB

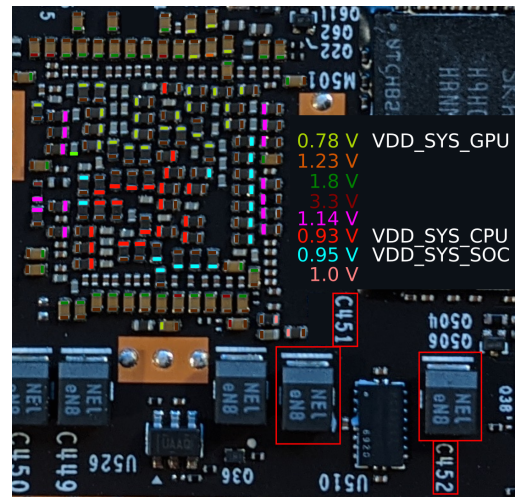


Figure 5: The different voltages on decoupling capacitors below the SoC.

below the SoC, see Fig. 5. In order to gain access to these capacitors during operation of the chip, we built an adapter using two SAMTEC connectors to place the module in an upside-down orientation, as shown in Fig. 4b.

The last question remaining is which power rail supplies the BPMP. To this regard, another observation can be used: The `VDD_SYS_SOC` rail is switched on first, while `VDD_SYS_CPU` and `VDD_SYS_GPU` are only enabled later during the boot process. Since the BPMP executes the first instructions, this leads to the conclusion that `VDD_SYS_SOC` is the rail supplying the BPMP.

At this point, we desoldered all decoupling capacitors shown in Fig. 5 that are connected to `VDD_SYS_SOC`. Furthermore, we removed the electrolytic capacitors with the PCB labels C451 and C452 (red boxes in the figure). Subsequently, we soldered the MOSFET used for injecting glitches to the PCB and connected its drain pin to one of the exposed pads – that was previously occupied by a decoupling capacitor – using a 35 mm piece of enameled copper wire with a diameter of 0.4 mm. This setup can also be seen in Fig. 4b. Afterward, we flashed a tightly coupled endless-loop as MB2, as described in Step 1 of the attack procedure. The program continuously emits the result of an addition operation on the UART interface after each iteration. Using a Python script to control the FPGA, we started glitching the X2 with increasing pulse lengths while manually triggering the glitch. We could observe corrupted data on the UART interface, confirming that we can enforce faulty behavior on the BPMP.

### B. Understanding the X2 Secure Boot

Following Step 2 of the attack procedure, we now needed to look for locations in the boot process that can potentially be exploited using FI. During our description of the X2 SoC, we stated that the Secure/TZ mode is not well described in the official documentation. In order to understand the necessary details, we relied on unofficial information that can be found online.

1) *Finding Hidden Documentation:* Nvidia offers header files containing memory mappings for the X2's registers as part of their L4T package. These mappings are also described in the X2's TRM [32], together with the information of which header describes which memory mapping. Since we were interested in the BR, we were looking for memory mappings that configure the iROM. The memory map lists BPMP\_ATCM as parent aperture for the BPMP\_BOOTROM, indicating that the BR is located within BPMP\_ATCM. When inspecting the memory mappings of the aperture, one mapping seems particularly interesting: BPMP\_ATCM\_CFG. The header file describing this mapping, `arbpmp_atcmcfg.h`, is missing in the L4T package and can only be found as part of Nvidia's Sensor Processing Engine source code [38]. When going through the file, some register names can be recognized for their similarity to important registers on the X1, like BPMP\_ATCMCFG\_SB\_CFG\_0 and BPMP\_ATCMCFG\_SB\_PiROM\_START\_0.

**Remark** Before realizing that the header file is offered by Nvidia, we searched for it on GitHub. Apart from finding a repository that includes the Nvidia code, the search also uncovered a repository called "switch-bootroms" [39]. This repository includes leaked BR source code for the Tegra SoCs with model numbers T210 and T214, whereas T210 is the original model of the Tegra X1 (codenamed "Erista"), and T214 is an updated version, also called Tegra X1+ (codenamed "Mariko") [40]. The X1+ includes faster clock speeds and, judging from comments and code in the repository, is hardened against FI. During our investigations, access to this code massively increased our understanding of the X2.

2) *ACCESS\_PiROM and PiROM\_START:* Looking at the header file mentioned above and going through the descriptions of the different bits in the named registers, we get an idea of how the protection mechanisms work. Three bits/registers stand out: SECURE\_BOOT, PiROM\_START, and ACCESS\_PiROM. Judging from the naming, default values, and the fact that similarly named bits exist on the X1, we can confidently assume their respective functionality. The PiROM\_START register holds an address marking the beginning of the protected iROM. No memory locations between PiROM\_START and the end of the iROM can be accessed unless ACCESS\_PiROM is set to 1. Both registers can only be changed while SECURE\_BOOT is set to 1. During MB2, this flag is set to 0.

Furthermore, a comment found in the source code from GitHub (line 492 of `nvboot_bpmp.c`) stresses the importance of preventing writes to PiROM\_START outside the Secure/TZ mode. Changing that value to the end of the iROM address range would allow an attacker to read the entire iROM content. Moreover, the comment confirms that the SECURE\_BOOT flag controls write access to PiROM\_START and ACCESS\_PiROM.

3) *Secrets in the BootROM:* For understanding the implications of leaking the BR, we consulted the official Nvidia documentation.

**MB1 Decryption Key** Looking at the X2's boot flow depicted in Fig. 2, we know that MB1 is decrypted by the BR stage. At this point, we can not know if BR accesses fuses to retrieve the key(s) or the protected iROM includes all necessary key material. Therefore, it is possible that MB1 can not be decrypted by only leaking the BR, but relevant fuse data or other protected content may also have to be leaked.

**Factory Secure Key Provisioning** Apart from key material to decrypt the second boot stage, different sources suggest the existence of a feature called Factory Secure Key Provisioning (FSKP). While no official documentation mentions this feature for the X2, Nvidia has a patent describing it [41]. The existence of the feature on the X2 is further supported by, e.g., a header file describing the respective keys [42]. Furthermore, we saw traces of the feature in the leaked BR code for the previous generation X1+ SoC.

The feature allows OEMs to encrypt data they want to burn into fuses. This is required when the OEM's threat model regards the factory – where the fuses are burnt – as potentially compromised. The encrypted data can only be decrypted using the FSKP key, that is provided by Nvidia and located in the X2's iROM. Since the keys are programmed into the devices during production, they can not be changed at a later point. There are 63 keys that Nvidia can assign, one per OEM. Since these keys are essentially the root of trust for an OEM's secure boot implementation, their secrecy is of utmost importance.

### C. Reverse Engineering unprotected BootRom

1) *Dumping iROM:* Once we gained a deeper understanding of the iROM protections using the files found on GitHub, we decided to read out the unprotected section of the iROM. Dumping the unprotected iROM is done by running a loop from the address marked as BOOTROM\_BASE in the TRM, until 128 kilobytes have been read or execution is interrupted by an exception due to unauthorized memory access. Inside the loop, each byte is sent via UART to the control PC.

Running this code in MB2 results in a dump of the address range `0x10000 - 0x11200`. The upper bound, `0x11200`, is the value stored in PiROM\_START during MB2. We could successfully import the read binary into Ghidra by selecting ARMv7, little-endian, 32-bit as processor type and entering the base address `0x10000` mentioned in the TRM. This allows us to analyze the BR code further.

2) *Analyzing iROM:* Execution on the Cortex-R5 starts with a `Reset` exception, forcing the program counter to address `0x0`. A branch to the reset handler is located at this address, which we call `reset()`. As a first step to analyzing `reset()`, we look up all memory addresses that Ghidra marks as unknown in the TRM. One accessed address of particular interest to us is BPMP\_ATCMCFG\_SB\_PiROM\_START\_0. We can see that the value of this address is updated from `0x400` to `0x2000`. This allows the BPMP to access the BR for subsequent initialization. While it may be possible to attack the store instruction with FI, it will most likely only break the boot process.

Apart from updating this address and the initialization of multiple processor status registers, two functions are called. We call the first function `ApplyIromPatches()`, which is responsible for activating a set of patches located in the X2's fuse memory. It enables Nvidia to update the BR in minor ways to fix critical bugs in the boot process. We call the second function `NonSecureDispatcher()`, where an array of function pointers is used to sequentially execute multiple functions to initialize the SoC before entering the protected address range. Lastly, the `reset()` function jumps into the protected iROM section at address `0x1200`.

Taking a closer look at the functions called by `NonSecureDispatcher()`, the second one stands out, as it looks like the one called `NvBootMainNonsecureRomEnter()` in the X1's source code. The function checks fuse data to determine whether or not the chip currently is in *Failure-AnalysisMode* or *PreproductionMode*. In case the chip is in *FailureAnalysisMode*, it executes a function called `NvBootUartDownload()`. This function is highly interesting for us, as it initializes a UART interface, sends out a prompt, and waits for data to arrive at the interface. Once the function has read enough data, it checks a simple checksum received together with the data. If the checksum is correct, the data is executed as code on the BPMP in the currently active privilege mode, i.e., the Secure/TZ mode. The prompt on the UART interface can serve as an indication of a successful glitch, as described in Step 2 of the attack procedure.

Since our X2 is booting normally, we can assume that neither of the above modes is active, and thus, the interface is not available to us. However, this code-loading feature is only protected by the fuse checks. There are no further security checks as to whether the loaded code is signed by an authorized party. Consequently, if we could manipulate the fuse check, we would likely gain code execution in the Secure/TZ mode.

3) *Fuse Check Code*: In order to understand whether or not the code can be glitched, we now take a closer look at the fuse check. The code shown in Listing 2 resembles the fuse check protecting the `NvBootUartDownload()` function. The two

```

1      push    {fp, lr}
2      bl     is_fam
3      cbz   r0, is_not_fam
4  is_fam_or_ppm:
5      bl     is_ppm
6      cbnz  r0, exit
7      bl     NvBootUartDownload
8  is_not_fam:
9      bl     is_ppm
10     cmp   r0, 0
11     bne   is_fam_or_ppm
12  exit:
13     pop   {fp, pc}

```

Listing 2: Pseudocode for fuse check protecting the `NvBootUartDownload()` routine.

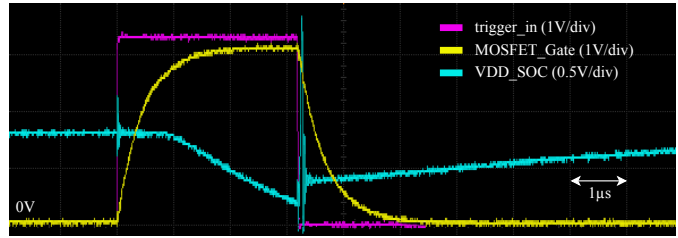


Figure 6: Example glitch of  $\approx 3\mu\text{s}$  length. `trigger_in` is the FPGA's GPIO signal controlling the multiplexer, `MOSFET_Gate` is measured at the MOSFET's gate, and `VDD_SOC` at the MOSFET's drain.

functions `is_fam` (*FailureAnalysisMode*) and `is_ppm` (*PreproductionMode*) always return 0 due to the values in the corresponding fuses. From line 3, the code normally jumps to line 8, `is_ppm` returns 0, and the instruction in line 11 does not branch. Therefore, the routine exits and the `NvBootUartDownload()` function is never reached.

However, if a glitch changes the branch direction in line 11 and the branch to `is_fam_or_ppm` is taken, the `NvBootUartDownload()` function is subsequently called, since `is_ppm` will still return 0 and the check in line 6 will not branch to `exit`. Similarly, if the branch in line 3 is skipped, we end up in line 5. Again, line 6 will not branch since `is_ppm` returns 0, and we also end up in `NvBootUartDownload()`. As a result of these observations, we hypothesize that the `NvBootUartDownload()` method can be activated by injecting a glitch at the correct time during boot.

#### D. Glitching the BootRom

1) *Proof of Concept*: To test the hypothesis of entering the UART bootloader by FI, we built an MB2 binary containing assembly code that structurally resembles the code shown in Listing 2. Instead of `is_ppm` and `is_fam`, the binary uses stubs that always return 0. The binary also includes a stable timing anchor in the form of a GPIO pin that is set to high before the critical code section begins. `NvBootUartDownload()` is replaced with a message sent over UART. After running some minor adjustments to the glitch parameters, we saw the success message being printed. This means that the code can indeed be glitched on the BPMP. Fig. 6 shows the oscilloscope trace of a single glitch.

Knowing that our hypothesis is correct, we will now describe how we attacked the actual boot process.

2) *Trigger Signal*: The trigger used for the proof of concept was a self-generated signal. However, for a reliable glitch on the actual target, we need a similar signal during booting, as described in Step 3 of the attack procedure. In order to find such a signal, we probed signals at different components on the PCB, like test points and voltage regulators, using an oscilloscope. However, none of the tested measuring points revealed trigger signals that potentially correlate with code executed in the early boot process. This is confirmed by further

analyzing the BR code, as no external signals are generated before the relevant code section starts.

Therefore, the only reliable signal at this point is the reset release signal. In many cases, the reset can not be used as a reliable trigger because there is jitter between reset release and the targeted code section, see Section IV-B. However, since we know from the analyzed BR code that there is only minimal code being executed between the reset exception and the fuse check, we can assume that the reset release works as a trigger. Therefore, we connect both the `trigger_in` and `reset_out` ports of the FPGA to a pin on the carrier board that exposes the module's reset signal `RESET_IN`, see Fig. 3. This way, we can pull `RESET_IN` low to initiate a full system reset and then trigger when we release the signal.

3) *Search Parameters:* After finding a usable trigger signal, we also need a sensible search space for the glitching parameters as described in Step 4 of the attack procedure. The parameters to optimize are glitch length and the offset from the trigger. On the one hand, we want to test as many glitch length/offset combinations as possible, but, on the other hand, including more parameters increases the time to search through all of them. In the following, we describe our approaches to minimize the search space for both parameters.

**Glitch Length** To optimize the glitch length as far as possible, we used our proof of concept MB2 binary. By having complete control over the trigger signal, we can run experiments in a very small offset range to find the glitch length interval that leads to the highest success rate. We marked the beginning and end of the critical code section by setting a GPIO pin high in the beginning and low after the section finished. Then we tried all relevant pulse lengths for every offset in that interval. First, we determined the range of plausible pulse lengths by setting the offset to 0 and then manually increasing the length of the pulse until the target reliably stops execution. This is the case for pulse lengths above 13  $\mu$ s. Subsequently, we tried every possible combination of both parameter ranges and recorded the parameter combination of successful glitches. After achieving ten successful glitches, we reduced the search intervals by constraining them to the minimum and maximum of the found parameters, plus some tolerance in either direction. We let the experiment run for eight hours and selected the three pulse lengths that worked most frequently, ending up with the following glitch pulse lengths: 11.30  $\mu$ s, 11.32  $\mu$ s, and 11.34  $\mu$ s.

**Offset Range** For limiting the interval to inject the glitch, we searched for signals and side-channels on the module that would allow us to identify the beginning of subsequent boot stages and operations. By our capability to execute code during MB2, we learned at which point the execution of MB2 starts. However, the time span between reset release and the start of MB2 is 172 ms. When searching the offset space with the highest possible resolution, one offset step is 20 ns long. Therefore there would be  $172 \text{ ms} / 20 \text{ ns} = 8600000$  possible offsets. Assuming that the result can be seen immediately after the glitch, one glitch attempt takes

$172 \text{ ms} / 2 = 86 \text{ ms}$  on average. This means that a full pass would take  $(8600000 * 86 \text{ ms}) / (1000 * 60 * 60 * 24) \approx 8.5$  days. Taking into account that there are three different pulse lengths that should be tested and, more importantly, that the success rate may be below 1 %, it seems problematic to use the start of MB2 as the limit for the offset. Therefore we needed to find an alternative way to minimize the offset interval that needs to be searched.

One crucial step for attacking the X1 is the ability to observe communication between the embedded MultiMediaCard (eMMC) and the X1 through a debug port [10]. The data observed on this port can be used as a reliable trigger signal. On the TX2 module, a comparable port is not available. Therefore, one avenue of research was to look for ways to load the MB1 image from another device than the eMMC, potentially allowing us to obtain a trigger signal during BR execution. We found a thread in the official Nvidia forum that discusses precisely this topic [43]. The original author describes how they use the strapping pins `BOOT_SELECT[2:0]` on their TX2i<sup>1</sup> to enter a mode where the BR starts looking for a Queued Serial Peripheral Interface (QSPI) device.

From our analysis of the BR in Section VI-C2, we know that the strapping pins are not checked before the fuse check. Furthermore, we know from the leaked source code that the function `SetupBootDevice()` in file `nvboot_coldboot.c` is responsible for checking the strapping pins. Since this function is located in the protected section of the BR, it will always be executed after the fuse check. Thus, we can further narrow down the search interval if we can trigger and identify the QSPI signals.

The OEM Product Design Guide [44] explains how the strapping pins are connected on the TX2 module. The pins are not exposed via the module connector but instead routed into a multiplexer located on the module, with pull-down resistors connected to the lines. We identified multiple multiplexer ICs on the module by their chip markings. By knowing to which pins of the module connector the multiplexer's output is routed and that there are three pull-down resistors nearby, we were able to identify the IC with the label U520 as the most likely candidate. By applying 3.3 V to the multiplexer inputs at reset release and measuring the signals `QSPI_SCK`, `QSPI_CS_N`, and `QSPI_IO0` (module connector names: G8, H8, and H7, respectively), we were able to observe the QSPI clock signal. This clock signal is only activated if the strapping pin is driven, starting 4.42 ms after reset release.

4) *Brute Force Search:* Since we know that the QSPI clock will only be activated after the fuse check, this gives us a much tighter upper bound to the offset interval. Applying the same calculation as earlier in Section VI-D3, we can approximate that a single run over the search space takes around 8.14 min. Since this seemed to be a reasonable time frame, we run the Python control code with the following parameters:

- **Pulse Lengths:** 11.30  $\mu$ s, 11.32  $\mu$ s, 11.34  $\mu$ s

<sup>1</sup>TX2 Industrial: Another Jetson module housing an X2 SoC with a different PCB and slightly different components

- **Offset Interval:** 0 ms - 4.42 ms with 20 ns steps

The success criterion was reading the text prompt mentioned in the BR analysis via UART. This happened twice, roughly eight hours after starting the experiment. The first two successful glitches were at offsets 2.633800ms and 2.625900ms with pulse lengths of 11.32 $\mu$ s and 11.34 $\mu$ s, respectively. To reproduce the result, we constrained the offset to an interval of 50 $\mu$ s around the earlier offset. In the resulting search interval between 2.585200ms and 2.634800ms and the same pulse lengths, the interface can be activated reliably within less than ten seconds. In most cases, it takes less than two seconds.

5) *Payload:* To leverage the debug interface's activation to gain code execution on the BPMP, the code that should be executed needs to be sent to the X2 over UART. This corresponds to Step 5 of the attack procedure. While there are no further security measures to authenticate this code, the payload has to adhere to the structure that the BR expects.

From the unprotected BR binary (see Section VI-C2), we can learn how to compose the payload. Cross-referencing the binary with the leaked code from GitHub indicates that the payload must start with a header called `NvBootUart_Header`. It includes the address of the first instruction, the length of the program following the header, and four unique identifications fields that are not relevant for us, as they are not checked. The actual program code follows the header. The last 4 bytes of the binary are a checksum composed by the one's complement of the sum of all preceding bytes.

Following this structure, we built an assembly program that adheres to the structure and sends the BR's protected region via UART to the control PC. That binary could then be saved and imported into Ghidra for further analysis.

#### E. Decrypting MB1

After dumping the full BR, we started analyzing the code that is responsible for loading and authenticating MB1. All cryptographic operations are executed by a hardware crypto engine. This engine also manages the keys used for these operations in several key slots. After analyzing the code that configures the hardware crypto engine, we could see that two keys called FEK are loaded into temporary key slots. In analogy to the naming in [38], we infer the name of these keys as Fuse Encryption Keys (FEKs). They are located in an aperture called `MISC` in the SoC. The first key loaded from these registers is named `FUSEROMENCRYPTIONNVKEY`. The second key is, depending on a fuse bit, loaded either from the same source or from registers called `FUSEROMENCRYPTIONTESTKEY` [38]. After the keys are loaded, a bit in an undocumented control register is set to protect these keys from being read again.

Next, several encrypted keys are read from the topmost 4 KiB of the BR memory region into other hardware key slots. These keys are then decrypted using Advanced Encryption Standard (AES) in Cipher Block Chaining (CBC) mode with the first FEK and an initialization vector (IV) of all zeros. The decrypted keys are internal to Nvidia and are used to decrypt assets loaded later in the boot chain. After that, the BR can

derive keys used to implement a secure boot mechanism for original design manufacturers (ODMs) using the second FEK, if this feature has been enabled by an ODM using a fuse bit.

The BR then proceeds to load the MB1 from non-volatile storage and to decrypt and authenticate it using the derived Nvidia MB1 keys. If an ODM wants to enable authentication/encryption for their software, i.e., for MB2 and later stages, MB1 is also encrypted and signed using the ODM's key material to protect the configuration of MB1 through the BCT. In such cases, MB1 is first authenticated/decrypted using the ODM key material, then using the Nvidia key material. Before jumping to the fully authenticated and decrypted MB1, the corresponding key slot containing the used keys for decryption is cleared to prevent it from being used by code loaded in later boot stages. All other key slots are kept intact and are used by MB1 to decrypt other boot loaders, metadata, and microcode.

Once we recognized that we were missing the FEKs to decrypt MB1, we enhanced the payload used to dump the BR to extract the registers containing the FEKs as well. The code path responsible for performing the UART pre-production serial boot that we entered using FI does not disable the readout of the FEKs. Therefore, we were able to dump the keys and use them together with the previously dumped BR to derive the MB1 decryption key, and decrypt the MB1 binary.

## VII. DISCUSSION

### A. Attack Applicability and Impact

Our work has shown that the boot processor on the Nvidia Tegra X2 SoC is susceptible to voltage FI. We uncovered a hidden UART bootloader and forced its activation by FI. This bootloader allows us to execute code in the Secure/TZ mode. Consequently, we could extract the full BR code, as well as all fuses and registers holding cryptographic keys. Using the recovered key material, we were able to decrypt the second-stage bootloader (MB1). Together with the extracted fuse values, an attacker would now be able to decrypt OEM's binaries of later boot stages. These capabilities form the perfect steppingstones to develop further attacks against a particular product.

Next to these insights, the UART bootloader would enable an adversary to take control over the boot process and the entire system. They could load their own unauthenticated code or disable security settings and then mimic the original system's behavior. Due to the short time of only a few seconds to a successful exploit, the FI can be conducted at each boot of the system without a significant increase in boot time. Moreover, in contrast to other more complicated FI techniques using lasers or EM radiation, this voltage FI capability can be persisted using a hardware implant. This has been demonstrated to be feasible, e.g., by Team Xecuter with their Nintendo Switch implant [45].

Our attack has shown that simple things like a bootloader that should only be used by Nvidia and is not officially documented can be leveraged to break the system's entire security. Devices and products of other manufacturers that are

not hardened against voltage FI will likely have similar weaknesses, which can be exploited easily with readily available equipment.

### B. Potential Mitigations

To protect SoCs against voltage FI attacks, we distinguish between mitigations that can be applied to prevent attacks on existing architectures and mitigations directly implemented into future chip generations.

1) *Protecting Existing SoCs*: After disclosing our findings to Nvidia, they proposed measures that can harden a device against voltage FI attacks. Since attackers need physical access for FI and similar attacks, OEM designers can increase the attack effort by adding board-level mitigations. For instance, heat- and solvent-resistant epoxy and non-removable metal shielding can prevent access to sensitive voltage rails and the removal of decoupling capacitors. Furthermore, embedding capacitive metal layers or decoupling capacitors into the PCB could reduce the effectiveness or accuracy of FI to a certain extent. Nevertheless, if the adversary only wants to exploit a single device, they might also consider developing a custom PCB for the SoC that facilitates access to the signal and omits decoupling capacitors.

2) *Protecting New SoC Versions*: The threat of fault injection for gaining adversarial code execution can be tackled from different directions. One could try to detect malicious voltage drops/glitches, and as a consequence, shut down the system to prevent further damage. Alternatively, one could try to prevent faulty execution in the presence of glitches, for instance, by introducing redundancy. These approaches imply changes in hardware or software design. Since neither approach can eradicate the chance of successful glitch injection completely, both approaches should be employed at the same time.

**Hardware-based Mitigations** Voltage monitoring circuits – as commonly implemented in modern smartcards – could help to detect glitches. Nvidia has recently patented a cross-domain voltage glitch detection circuit, which can be implemented into an SoC [46]. The main idea is that an SoC contains multiple voltage rails. Therefore, circuits in independent voltage domains can monitor voltage levels in other domains, and if there is a glitch on a specific rail, assert an alert signal. Although we think this is a promising approach, it should be kept in mind that there might exist voltage glitch shapes that can cause faulty behavior but can not be detected by a particular protection circuit.

Fully integrating the voltage regulators into the SoC could be another solution. However, faults can not only be induced by glitching the supply voltage. In the past couple of years, EM fault injection techniques against modern CPUs have been examined to inject faults in a targeted and contactless way [8], [47]. Consequently, a holistic view is necessary to prevent all kinds of fault injection attacks that can manipulate the behavior of the target device.

**Software-based Mitigations** Hardening the BR might be another option to prevent the adversary from gaining code

execution. However, this is a complex task since the characteristics and potentials of faults are not well understood. Particularly, there is no model which covers all possible faults. Nevertheless, Riscure and ARM propose countermeasures that can decrease the probability of successful attacks [48], [49]. For instance, constants with large hamming distances can complicate flipping one valid value to another, and double-checks can protect branch conditions. Moreover, loop integrity checks make sure that the loop exits as intended, and a global counter can be used to monitor the program flow and detect anomalies. For assessing software countermeasures against fault attacks, different simulation-based frameworks have been proposed [50], [51].

It should be noted that Nvidia is seemingly aware of these countermeasures. We saw most of these countermeasures being used the leaked X1+ "Mariko" BR source code. Unfortunately, they are not employed in the X2. The general approach of software-based mitigations might be promising as they can also protect against fault attacks other than voltage glitching.

## VIII. CONCLUSION

Starting with the knowledge that the Nvidia Tegra X1 SoC is susceptible to voltage FI, we have investigated the boot security on its successor, the Tegra X2. We have shown that a hidden bootloader, which allows code execution with the highest privileges, can be enabled using voltage FI. Using this capability, we were able to extract the entire bootloader's code and decryption keys for later boot stages. Furthermore, we described that an adversary could use the bootloader for injecting unauthentic code to bypass trusted code execution on the system. After disclosing our findings to Nvidia, they confirmed the vulnerability and stated that future device generations contain mitigations against FI attacks. As devices of other manufacturers potentially are vulnerable in a similar manner, we have discussed countermeasures to mitigate our attack. Manufacturers and designers should not forget about seemingly simple hardware attacks that have been around for already more than two decades.

## ACKNOWLEDGMENT

The work described in this paper has been supported by the Einstein Foundation in form of an Einstein professorship (EP-2018-480).

## REFERENCES

- [1] H. Bar-El, H. Choukri, D. Naccache, M. Tunstall, and C. Whelan, "The Sorcerer's Apprentice Guide to Fault Attacks," *Proceedings of the IEEE*, vol. 94, no. 2, pp. 370–382, Feb. 2006.
- [2] O. Koemmerling and M. G. Kuhn, "Design Principles for Tamper-Resistant Smartcard Processors," *Smartcard 99*, pp. 9–20, 1999.
- [3] D. Boneh, R. A. DeMillo, and R. J. Lipton, "On the Importance of Eliminating Errors in Cryptographic Computations," *Journal of Cryptology*, vol. 14, no. 2, pp. 101–119, Mar. 2001.
- [4] N. Timmers, A. Spruyt, and M. Witteman, "Controlling PC on ARM Using Fault Injection," in *2016 Workshop on Fault Diagnosis and Tolerance in Cryptography (FDTC)*, Aug. 2016, pp. 25–35.
- [5] Y. Lu. (2019) Injecting Software Vulnerabilities with Voltage Glitching. [Online]. Available: <https://arxiv.org/abs/1903.08102>

- [6] A. Vasselle, H. Thiebeault, Q. Maouhoub, A. Morisset, and S. Ermeneux, "Laser-Induced Fault Injection on Smartphone Bypassing the Secure Boot-Extended Version," *IEEE Transactions on Computers*, vol. 69, no. 10, pp. 1449–1459, Oct. 2020.
- [7] K. M. Abdellatif and O. Hériveaux, "SiliconToaster: A Cheap and Programmable EM Injector for Extracting Secrets," in *2020 Workshop on Fault Detection and Tolerance in Cryptography (FDTC)*, Sep. 2020, pp. 35–40.
- [8] T. Trouchkine, S. K. Bukasa, M. Escouteloup, R. Lashermes, and G. Bouffard, "Electromagnetic Fault Injection against a Complex CPU, toward New Micro-Architectural Fault Models," *Journal of Cryptographic Engineering*, Mar. 2021.
- [9] A. Djellid-Ouar, G. Cathebras, and F. Bancel, "Supply voltage glitches effects on CMOS circuits," in *International Conference on Design and Test of Integrated Systems in Nanoscale Technology, 2006. DTIS 2006. IEEE*, 2006, pp. 257–261.
- [10] A. Galauner. (2018) Glitching the Switch. Video. [Online]. Available: <https://media.ccc.de/v/c4.openchaos.2018.06.glitching-the-switch>
- [11] Nvidia Corp. (2021) NVIDIA DRIVE PX - Scalable AI Supercomputer For Autonomous Driving. [Online]. Available: [https://www.nvidia.com/content/nvidiaGDC/sg/en\\_SG/self-driving-cars/drive-px/](https://www.nvidia.com/content/nvidiaGDC/sg/en_SG/self-driving-cars/drive-px/)
- [12] —. (2016) Tesla Motors' Self-Driving Car "Supercomputer" Powered by NVIDIA DRIVE PX 2 Technology. [Online]. Available: <https://blogs.nvidia.com/blog/2016/10/20/tesla-motors-self-driving/>
- [13] Daimler AG. (2018) At a Glance: The Key Data on MBUX. [Online]. Available: <https://media.daimler.com/marsMediaSite/ko/en/32705799>
- [14] S. Abuelsamid. (2020) Hyundai Adopts Nvidia Parker 'System On A Chip' For Infotainment. [Online]. Available: <https://www.forbes.com/sites/samabuelsamid/2020/11/11/hyundai-motor-group-adopts-nvidia-parker-soc-for-infotainment/>
- [15] C. O'Flynn and Z. D. Chen, "ChipWhisperer: An Open-Source Platform for Hardware Embedded Security Research," in *Constructive Side-Channel Analysis and Secure Design*. Springer International Publishing, 2014, pp. 243–260.
- [16] C. Bozzato, R. Focardi, and F. Palmari, "Shaping the Glitch: Optimizing Voltage Fault Injection Attacks," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 199–224, Feb. 2019.
- [17] 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," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021.
- [18] K. Murdock, D. Oswald, F. D. Garcia, J. V. Bulck, F. Piessens, and D. Gruss, "Plundervolt: How a Little Bit of Undervolting Can Create a Lot of Trouble," *IEEE Security & Privacy*, vol. 18, no. 5, pp. 28–37, Sep. 2020.
- [19] P. Qiu, D. Wang, Y. Lyu, R. Tian, C. Wang, and G. Qu, "VoltJockey: A New Dynamic Voltage Scaling based Fault Injection Attack on Intel SGX," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–1, 2020.
- [20] C. O'Flynn, "Fault Injection using Crowbars on Embedded Systems," *IACR Cryptol. ePrint Arch.*, 2016. [Online]. Available: <https://eprint.iacr.org/2016/810>
- [21] A. Milburn, N. Timmers, N. Wiersma, R. Pareja, and S. Cordoba, "There Will Be Glitches: Extracting and Analyzing Automotive Firmware Efficiently," *Black Hat USA*, 2018.
- [22] J. V. den Herrewegen, D. Oswald, F. D. Garcia, and Q. Temeiza, "Fill your Boots: Enhanced Embedded Bootloader Exploits via Fault Injection and Binary Analysis," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 56–81, 2021.
- [23] Nate Lawson. How the PS3 Hypervisor was Hacked. [Online]. Available: <https://rdist.root.org/2010/01/27/how-the-ps3-hypervisor-was-hacked/>
- [24] Razkar. The Reset Glitch Hack - A New Exploit on Xbox 360. [Online]. Available: <http://www.logic-sunrise.com/news-341321-the-reset-glitch-hack-a-new-exploit-on-xbox-360-en.html>
- [25] TechInsights Inc. Nintendo Switch Teardown. [Online]. Available: <https://www.techinsights.com/blog/nintendo-switch-teardown>
- [26] M. Szekely. (2021, Jul.) Qyriad/fusee-launcher. [Online]. Available: [https://github.com/Qyriad/fusee-launcher/blob/master/report/fusee\\_gelee.md](https://github.com/Qyriad/fusee-launcher/blob/master/report/fusee_gelee.md)
- [27] G. T. H. G. I. Roussel-Tarbouriech, N. Menard, T. True, T. Vi, and Reisyukaku. (2019) Methodically Defeating Nintendo Switch Security. [Online]. Available: <https://arxiv.org/abs/1905.07643>
- [28] plutoo, derrek, and naehrwert. (2017) Console security Switch. Video. [Online]. Available: [https://media.ccc.de/v/34c3-8941-console\\_security\\_-\\_switch](https://media.ccc.de/v/34c3-8941-console_security_-_switch)
- [29] Nvidia Corp. (2017) Jetson TX2 Module. [Online]. Available: <https://developer.nvidia.com/embedded/jetson-tx2>
- [30] —. (2017) Harness AI at the Edge with the Jetson TX2 Developer Kit. [Online]. Available: <https://developer.nvidia.com/embedded/jetson-tx2-developer-kit>
- [31] —. Jetson TX2 Boot Flow. [Online]. Available: [https://docs.nvidia.com/jetson/l4t/index.html#page/Tegra%20Linux%20Driver%20Package%20Development%20Guide/bootflow\\_tx2.html](https://docs.nvidia.com/jetson/l4t/index.html#page/Tegra%20Linux%20Driver%20Package%20Development%20Guide/bootflow_tx2.html)
- [32] —, *Technical Reference Manual Nvidia Parker Series SoC, v1.0p*, Jun. 2017.
- [33] Thomas Roth. Chipfail-glitcher. GitHub. [Online]. Available: <https://github.com/chipfail/chipfail-glitcher>
- [34] S. Gayou. (2020) Sgayou/rbasefind. GitHub. [Online]. Available: <https://github.com/sgayou/rbasefind>
- [35] Nvidia Corp., *OpenVReg: Open Voltage Regulator Specification*, 2011. [Online]. Available: <https://international.download.nvidia.com/openvreg/openvreg-type2-plus-1-specification.pdf>
- [36] —, *Linux For Tegra*, 32nd ed., Nvidia Corporation, Feb. 2021. [Online]. Available: <https://developer.nvidia.com/embedded/linux-tegra>
- [37] —, *Tegra Linux Driver Package Development Guide 27.1*, Mar. 2017. [Online]. Available: <https://docs.nvidia.com/jetson/archives/l4t-archived/l4t-271/index.html>
- [38] —. Jetson Sensor Processing Engine (SPE). Source code. [Online]. Available: [https://developer.nvidia.com/embedded/L4T/r32\\_Release\\_v5.0/sources/T186/l4t\\_rt\\_aux\\_cpu\\_src.tbz2](https://developer.nvidia.com/embedded/L4T/r32_Release_v5.0/sources/T186/l4t_rt_aux_cpu_src.tbz2)
- [39] highmirai. (2021) Zerospace-nx/switch-bootroms. GitHub. [Online]. Available: <https://github.com/zerospace-nx/switch-bootroms>
- [40] Wikipedia. (2021) Tegra. [Online]. Available: [https://en.wikipedia.org/wiki/Tegra#Tegra\\_X1](https://en.wikipedia.org/wiki/Tegra#Tegra_X1)
- [41] J. Huang, P. Chou, and A. WOO, "Secure provisioning of semiconductor chips in untrusted manufacturing factories," US Patent US10 387 653B2, 2019. [Online]. Available: <https://patents.google.com/patent/US10387653B2/en>
- [42] R. Abdel-Tawab. (2019) Android\_bootable\_bootloader\_cboot-t186. GitHub. [Online]. Available: [https://github.com/Rashed97/android\\_bootable\\_bootloader\\_cboot-t186/blob/lineage-16.0/hwinc-t18x/nvboot\\_se\\_aes.h#L145-L221](https://github.com/Rashed97/android_bootable_bootloader_cboot-t186/blob/lineage-16.0/hwinc-t18x/nvboot_se_aes.h#L145-L221)
- [43] Nvidia Corp. (2019) Boot source other than eMMC or USB. Forum entry. [Online]. Available: <https://forums.developer.nvidia.com/t/boot-source-other-than-emmc-or-usb/81621>
- [44] —, *OEM Product Design Guide Nvidia Jetson TX2 Series*, May 2020.
- [45] T. Xecuter. (2021) SX Core Manual. [Online]. Available: [https://sx.xecuter.com/download/manuals/sxcore/\[EN\]\\_SX\\_Core\\_v1.1.pdf](https://sx.xecuter.com/download/manuals/sxcore/[EN]_SX_Core_v1.1.pdf)
- [46] K. Rajpathak and T. Raja, "Cross domain voltage glitch detection circuit for enhancing chip security," US Patent US20 200 285 780A1, Sep., 2020. [Online]. Available: <https://patents.google.com/patent/US20200285780A1/en>
- [47] T. Trouchkine, G. Bouffard, and J. Clédière, "Fault Injection Characterization on Modern CPUs," in *Information Security Theory and Practice*, ser. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2020, pp. 123–138.
- [48] M. Witteman. (2018) Riscure: Secure Application Programming in the Presence of Side Channel Attacks. [Online]. Available: [https://www.riscure.com/uploads/2018/11/201708\\_Riscure\\_Whitepaper\\_Side\\_Channel\\_Patterns.pdf](https://www.riscure.com/uploads/2018/11/201708_Riscure_Whitepaper_Side_Channel_Patterns.pdf)
- [49] T. Ban. (2020) Arm Ltd.: Trusted Firmware M. Presentation slides. [Online]. Available: [https://www.trustedfirmware.org/docs/TF-M\\_fault\\_injection\\_mitigation.pdf](https://www.trustedfirmware.org/docs/TF-M_fault_injection_mitigation.pdf)
- [50] A. Höller, A. Krieg, T. Rauter, J. Iber, and C. Kreiner, "QEMU-Based Fault Injection for a System-Level Analysis of Software Countermeasures Against Fault Attacks," in *2015 Euromicro Conference on Digital System Design*, Aug. 2015, pp. 530–533.
- [51] H. Schirmeier, M. Hoffmann, C. Dietrich, M. Lenz, D. Lohmann, and O. Spinczyk, "FAIL\*: An Open and Versatile Fault-Injection Framework for the Assessment of Software-Implemented Hardware Fault Tolerance," in *2015 11th European Dependable Computing Conference (EDCC)*. IEEE, Sep. 2015, pp. 245–255.