High Fidelity Security Mesh Monitoring using Low-Cost, Embedded Time Domain Reflectometry

Jan Sebastian Götte[1] and Björn Scheuermann[2]

[1] Technical University of Darmstadt, Darmstadt, Germany, research@jaseg.de
[2] Technical University of Darmstadt, Darmstadt, Germany,
bjoern.scheuermann@kom.tu-darmstadt.de

Abstract. Security Meshes are patterns of sensing traces covering an area that are used in Hardware Security Modules (HSMs) and other systems to detect attempts to physically intrude into the device's protective shell. State-of-the-art solutions manufacture meshes in bespoke processes from carefully chosen materials, which is expensive and makes replication challenging. Additionally, state-of-the-art monitoring circuits sacrifice either monitoring precision or cost efficiency. In this paper, we present an embeddable security mesh monitoring circuit constructed from low-cost, standard components that utilizes Time Domain Reflectometry (TDR) to create a unique fingerprint of a mesh. Our approach is both low-cost and precise, and enables the use of inexpensive standard Printed Circuit Boards (PCBs) as security mesh material. We demonstrate a working prototype of our TDR circuit costing less than 10 € in components that achieves both time resolution and rise time better than 200 ps—a 25× improvement over previous work. We demonstrate a simple classifier that detects several types of advanced attacks such as probing using an oscilloscope probe or micro-soldering attacks with no false negatives.

Keywords: Tamper Sensing · Tamper Response · Physical Security · Security Mesh · Hardware Security Module (HSM) · FIPS 140-2/3 · ISO/IEC 24759 · PCI PTS HSM MSR

1 Introduction

Security meshes continue to be the state of the art for tamper sensing in applications where sophisticated physical attacks such as attempts at drilling or sawing through the device's enclosure to place probes must be prevented. Common applications for such meshes include Hardware Security Modules (HSMs) used to store and process cryptographic keys applying security standards such as FIPS-140-2 [1] or ISO/IEC 24759 [18]. Other applications include card payment terminals where PCI PTS HSM standards [32] are applicable. Security meshes usually consist of two or more conductive traces that are laid out in a meandering pattern to cover a surface. A sensing circuit electrically monitors these traces to detect attempts at penetrating this surface.

As is often the case with security technologies, in practice a tension exists between the level of security offered by a particular security mesh implementation and its implementation cost. Commercial designs often only coarsely monitor the conductivity of the mesh traces and are incapable of detecting attacks that manipulate small parts of the mesh. The most secure meshes are made in custom manufacturing processes. Materials such as polymer substrates are specifically chosen such that the mesh is difficult to manipulate without breaking it. A drawback of this approach is that the specialized manufacturing

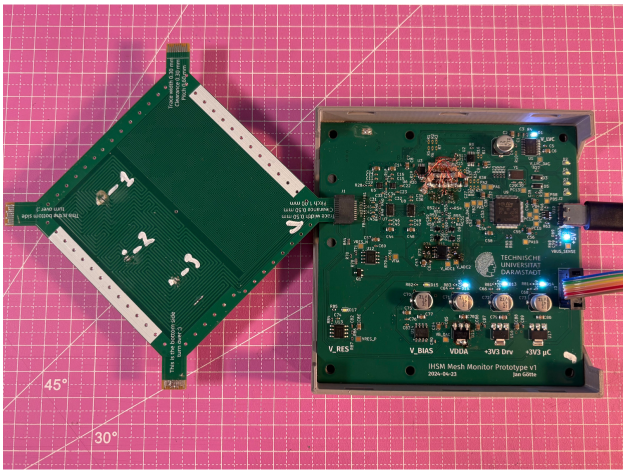

Figure 1: Measurement setup. Shown are the test specimen board on the left, and the frontend board with one of the four pulse amplifiers in the center. The frontend board is powered through a USB-C connection, and data is sent to a computer through a Single-Wire Debug (SWD) interface. The grid in the background has 10 mm pitch.

processes are difficult to replicate and that the resulting cost of the mesh is high. In some lower-security applications such as card payment terminals, simpler approaches are still commonly used for their ease of implementation. Often, standard copper/polyimide Flexible Printed Circuits (FPCs) or even standard Printed Circuit Boards (PCBs) are used because of the wide availability of manufacturing services.

Several academic approaches exist that target low-cost [47, 45, 9, 46] or high-performance mesh monitoring [16, 17, 12]. Some academic works even try to replace the security mesh with entirely different tamper sensing primitives [40, 44]. High-performance mesh monitoring approaches try to characterize the mesh's physical properties with high accuracy, but often come at the cost of specialized, expensive circuitry. Low-cost approaches utilize advanced analog techniques in their circuitry to extract precise measurements using few components. They trade off measurement precision for lower component cost. Besides simple monitoring, detecting tamper attempts by replacing the mesh with a macro-scale Physically Unclonable Function (PUF) has also been researched [16, 40, 44], albeit this comes with complex monitoring circuits that utilize expensive, specialty components.

To enable the use of less expensive, commodity materials such as Printed Circuit Boards (PCBs) without compromising security, mesh integrity must be monitored with high fidelity. In this paper, we present a low-cost monitoring circuit for security meshes that combines Time Domain Reflectometry (TDR) with equivalent time sampling. Our approach provides high measurement fidelity and enables the use of meshes made from less expensive materials in high-security applications.

Our circuit generates a very fast pulse with a rise time lower than 200 ps that is broadcast into the mesh. While the pulse traverses the mesh, parts of its energy are reflected on imperfections inside the mesh, including those caused by tampering attempts. Our circuit uses a fast, low-cost equivalent time sampling frontend to receive, amplify and record these reflections to create a *fingerprint* of the mesh that is highly sensitive to changes caused by tampering.

We demonstrate a working prototype of our design and present practical measurements of its electrical parameters as well as its performance under several practical attack scenarios. A photo of our prototype setup including a security mesh specimen is shown in Figure 1.

Compared to previous academic designs, our approach can be implemented at a lower cost using exclusively inexpensive, commercially available mass-market components. Our TDR frontend improves upon previous, delay-based approaches in monitoring fidelity [47, 45]. Our design achieves sufficient sensitivity to detect high-impedance oscilloscope probes despite such probes being specifically designed to conduct measurements without disturbing the circuit under test. Unlike previous, capacitance-based approaches, our design is compatible with inexpensive signal switch ICs, enabling the protection of arbitrarily large meshes at minimal cost without compromising sensitivity.

The contributions of our work are as follows:

- To our knowledge, our design is the first to apply a low-cost embedded differential Time Domain Reflectometry (TDR) frontend to security mesh monitoring. Our design achieves pulse rise times below 200 ps, a 25× improvement over the closest previous work [47, 45].

- Our approach provides higher fidelity compared to state-of-the-art security mesh conductivity monitoring or previous low-cost approaches. It enables the use of meshes manufactured using less advanced technologies such as standard FPC or PCB processes. Our TDR frontend produces 70 data points for each meter of mesh length, resulting in a measurement density per mesh area of 200 bit/cm^2 when using a 200 µm pitch mesh manufactured in a standard low-cost PCB process.

- We present a working prototype along with extensive experimental results, including laboratory performance measurements. We practically demonstrate that our design is able to not only detect but distinguish and even localize attacks in several realistic attack scenarios.

- Our design is based entirely on commercially available, inexpensive mass-market components. It can be replicated and improved without access to bespoke production equipment or semiconductor manufacturing capabilities. To facilitate further research and practical applications, we publish our prototype under an Open Source license.

2 Related Work

Tamper sensing meshes are used in numerous applications from Hardware Security Modules (HSMs) to card payment terminals [2, 41]. Despite their widespread use, security mesh design and monitoring is covered by a sparse research corpus. Commercially, security-by-obscurity is often considered a good idea and little detail is published on physical security implementations [3].

Patent literature gives a partial view of commercial developments in this area. Even in recent patents such as [6, 29, 34, 48, 21, 8] from HSM manufacturers IBM and HP, ATM component manufacturer Cryptera, payment terminal manufacturer Stripe, and chip manufacturers Texas Instruments and Zilog, cited monitoring methods are basic and do not go beyond a simple measurement of resistance or capacitance.

Academic research in the area is more advanced and spans both improvements to security meshes and their monitoring circuits [16, 9, 46], as well as approaches that entirely replace the security mesh with other primitives based on e.g. radio frequency or optical measurements that aim to sense tampering with a device [40, 44]. A drawback of techniques aiming to replace security meshes with other sensor types is that it is difficult to prove such sensors do not have blind spots.

2.1 Security Mesh Monitoring and Design

Meshes as capacitive PUFs. Immler et al. [16], Obermaier et al. [30], and Garb [12] propose one of the most advanced security mesh designs in the current academic state of the art. They use a specialized security mesh as a Physically Unclonable Function (PUF), combining tamper sensing with cryptographic key storage. In their design, the mesh consists of a cross-hatch pattern made from several dozen individually addressable capacitive electrodes. They manufacture their meshes in a specialized process that results in unpredictable, random variations in capacitance between electrodes. They propose an analog frontend that measures the precise mutual capacitance of each pair of electrodes [30] using an approach similar to Sato, Poupyrev, and Harrison [38], and they use the resulting capacitance matrix as the basis of their PUF. In further work, they demonstrate a custom IC integrating the monitoring circuit [11].

Advantages of their system include high sensitivity to modifications, as well as that as a PUF, the system does not require a continuous power supply. Disadvantages include the limited mesh size a single circuit can support due to dynamic range constraints, the specialized manufacturing process needed for the mesh as well as the high cost of the monitoring circuit. Common physical security standards require systems to actively destroy all key material when tampering is detected [1, 18, 32]. Like other PUF-based systems, their system naturally lacks this capability.

Key differences of our system include:

- Our system can cover larger meshes without loss of precision using a single TDR frontend through multiplexing.
- Our system supports meshes manufactured using standard, low-cost PCB processes.
- Our design requires only widely available, low-cost commodity components, for each of which alternatives from other manufacturers are available.
- Our approach has improved resiliency to electromagnetic interference and works with unshielded meshes.

Bridge measurement of capacitive interdigital meshes. Dupont et al. [9] introduce a simple analog circuit approach for monitoring meshes laid out as a set of capacitive interdigital structures not unlike the combs found in Micro-Electromechanical System (MEMS) accelerometers and gyroscopes. They subdivide the mesh into four equal-size quadrants, each containing two equal-size interdigital electrodes. They connect the resulting eight electrodes in a capacitive bridge configuration and measure the bridge's balance using a simple analog monitoring circuit based on homodyne detection. Advantages of their system include the simple, low-power monitoring circuit made from basic, cheap components and the capability to work with single-layer meshes such as those produced using Laser Direct Structuring (LDS). From a security point of view, a drawback of their approach is that to achieve its low-power usage, measurement resolution is sacrificed and all information on the mesh's state is collapsed into a single, scalar measurement.

Frequency-domain mesh characterization. Vasile and Svasta [46] introduce a monitoring method where they feed a variable-frequency signal into one end of a continuous mesh trace, and measure the power of the signal coming out of the other end. In essence, their setup measures S_{12} magnitude in a similar way to a network analyzer.

Advantages of their design include the simple implementation and the potentially robust nature of frequency-domain measurements. Disadvantages include a nonstandard three-layer mesh stackup, as well as the susceptibility of the system to attack by emulation given that the log power sensor they are using at the mesh output is designed to be insensitive to any signal characteristics apart from total signal power.

Time domain mesh monitoring. Time-Domain Reflectometry has been proposed for tamper sensing in nuclear arms control applications [31]. However, compared to our design, the systems proposed in this field are usually much larger, using standard benchtop measurement equipment to perform TDR. Additionally, they target lower time resolution since they are designed to monitor spans of cable up to several hundred meters in length.

Closest to our proposal in the academic corpus is the work of Vasile et al. [47] and Vasile and Svasta [45], where they propose monitoring the time domain response of a mesh using a circuit made from a pulse generator and a fast Analog-to-Digital Converter (ADC). To avoid an expensive, high-speed digital processing pipeline, their design is centered around a specialized high-speed ADC that has a built-in sample memory. Using this part, they capture a pulse at high speed after it traverses the mesh. Subsequently, they slowly process the captured data from memory.

Advantages of their design include better sensitivity to changes in total mesh trace length compared to simple continuity monitoring and the low complexity of their analog frontend. Disadvantages include the reliance on a specialty ADC that cannot easily be replaced with any other commercially available component and the coarse time resolution.

Key differences between their design and our proposal include:

- Their design is sensitive to total length, but not to the location of faults. Their design measures the mesh's *transmission* characteristic, which collapses detail about faults along the mesh into a small number of ADC samples at the pulse edge. Using such a measurement, it is not possible to localize faults. In contrast, our approach measures the signal's *reflected* component, which spreads information over time and enables us to localize faults.

- Our design uses only inexpensive, widely available parts. All parts in our design can easily be substituted for other, similar parts from different manufacturers.

- Our approach provides 25× higher time resolution through Equivalent Time Sampling. This is a fundamental limitation of their design, as the cost of ADCs and their associated circuitry increases steeply with speed[1].

2.2 Equivalent Time Sampling

Today, systems that digitize high-speed signals usually use a fast ADC, sometimes preceded by one or several downconverting mixers. This development was enabled by both the increasing availability of ADCs capable of digitizing hundreds of megasamples per second at a reasonable resolution, and by the increase in speed of CPUs, FPGAs, and other components of the digital processing chain. However, this is largely a development of this millennium–meanwhile, signals far into the gigahertz range have been studied since the advent of radar technology in the Second World War [19]. Enabled by the progress from vacuum tubes to semiconductor devices, equivalent time sampling became the technology of choice for the latter half of the twentieth century until around the turn of the millennium the introduction of high-speed digital processing and fast ADCs enabled real-time conversion up into higher microwave frequencies, today reaching beyond the 100 GHz boundary.

Kahrs [19] trace back the style of four-diode balanced bridge sampling gate that we use to a vacuum tube implementation presented in Chance et al. [7]. This style of sampling gate found application in a number of sampling oscilloscopes throughout the twentieth century in several oscilloscope sampling frontends such as HP's 187B [15].

While initially equivalent time sampling was used to circumvent technological limitations, more recently it has also been used to achieve cost-optimized designs [14]. Going

[1] For reference, the least expensive ADC available at distributor DigiKey that would match the 200 ps time resolution of our approach would cost 320 € at quantity 100 and require national security clearance for export from its manufacturer in the USA.

along similar principles, Polášek [33] presents a design for a minimal sampling TDR circuit that uses a CMOS clock generator IC along with a CML fanout buffer for pulse generation. The circuit improves upon the double sampling design first presented by Houtman [14] to reconstruct a downsampled copy of the input signal in the analog domain before digitization.

2.3 Low-Cost Time Domain Reflectometry

Bencivenni et al. [4] present an FPGA-based embedded reflectometer design. Since their design is based on an early FPGA family dating back to 2003 that lacked the speed and the adjustable I/O delay features of more modern FPGA families, their design uses the FPGA's logic resources to achieve adjustable delays. Negrea and Rangu [28] show an equivalent time sampling TDR that uses specialized adjustable delay line ICs for pulse generation. Lee, Sung, and Park [20] achieve very high time resolution in an equivalent time sampling TDR system by using a vernier approach to pulse generation, such that their system is limited by analog bandwidth, not time resolution. Trebbels et al. [43] show another FPGA-based TDR. Their system also uses a part from the same early FPGA family as Bencivenni et al. [4], and they work around its lack of precise timing primitives by generating a low-frequency sine wave through DDS, which they filter, and then sample using a comparator - a similar approach to the timing generation in Houtman [14]. Additionally, they avoid the need for a discrete ADC by implementing a $\Delta\Sigma$ loop around a fast comparator, trading off slower acquisition time for lower hardware complexity. They use a $5.5\,\mathrm{V\,ns^{-1}}$ slew rate wideband amplifier IC to generate their stimulus pulse, achieving a rise time of 2 ns. As a result, similar to Lee, Sung, and Park [20], their design is limited by analog bandwidth–here resulting from the nanosecond-scale stimulus rise time–not by frontend time resolution. Compared with this and other previous approaches, our proposed system is not only faster, but presents a more balanced trade-off between time resolution and analog bandwidth.

2.4 Device Fingerprinting through Impedance Sensing

Recently, impedance analysis on the Power Distribution Network (PDN) of PCB assemblies has been proposed as a fingerprinting technique aimed at detecting Hardware Trojans (HT) inserted into a board [10, 24]. Usually, all chips on a board are directly connected to the board's PDN. Thus, characterizing the board's PDN does not only yield information on possible modifications to the board's PDN itself—such as modified traces or removed passive components—it also reflects information about the internal structure of chips connected to the PDN. Impedance analysis techniques generally probe the circuit during operation using high-frequency signals. They have been proven using an external Vector Network Analyzer in one-Port [26] configuration measuring reflected signal components as well as using two or more ports measuring transmitted signal components [50]. Both Time Domain Reflectometry [10] and conventional frequency-domain VNA measurements [24] have been shown to be effective. From a signal theory point of view, both techniques can be considered equivalent.

While using an external VNA is feasible for validation in a factory setting, several research works embed the measuring system into the PCB as either a discrete circuit [10] or as part of an FPGA gateware [24, 25]. With such a system, boards can self-verify in the field after deployment, enabling the use of the system for active tamper sensing. While at less than 2 GHz the achievable bandwith of such systems is lower than that provided by an external, research-grade VNA, it turns out that the frequencies of interest in the impedance profile of practical boards lie inside of this small bandwidth [24].

Variations of impedance analysis techniques have been demonstrated that detect changes inside individual chips using board-level measurements [22], that detect manipulatoins using non-contact near-field Radio Frequency (RF) measurements [36], that detect the mechanical preparation of a target chip for backside attacks using onboard measurements [25], and that adapt the technique as an offensive tool for side-channel analysis (SCA) attacks [23].

Similar to PDN impedance analysis, our proposed technique also embeds a RF measurement circuit in a target board. TDR and frequency-domain VNA measurements resolve the same information about a target circuit from a signal theory perspective. Our system reaches a significantly higher bandwidth than embedded measurement setups from differs from PDN impedance analysis literature, and that our proposed tamper-sensing meshes are specifically built as sensors. Our technique is better suited to active tamper-sensing applications where the sensing circuit is continuously powered. In contrast to PDN impedance analysis techniques that need the entire PDN to be powered, our proposed technique can be applied to protect an unpowered payload circuit. In a practical application, both PDN impedance analysis and TDR-based tamper-sensing meshes could complement each other to form a comprehensive defense where PDN impedance analysis checks the core system's integrity, with TDR-based meshes covering everything outside the purview of PDN impedance analysis.

3 Monitoring a Security Mesh using Time Domain Reflectometry

Time Domain Reflectometry (TDR) is a well-known technique that is used to locate faults along a signal channel such as a copper cable, or an optical fiber. In TDR, a pulse is sent into the beginning of the channel. While the pulse traverses the channel, any fault such as a discontinuity in electrical impedance or optical density causes part of the pulse to travel back in a partial reflection. TDR monitors these reflections returning to the beginning of the channel by recording the signal measured at it after the pulse has been sent. When the pulse reaches the end of the channel, depending on termination it can be reflected to travel back to the beginning, which allows measurement of the channel's length.

3.1 Attacks on a Security Mesh Viewed Using TDR

In this paper, we apply TDR to monitor a security mesh for changes caused by an attack. Our prototype setup consists of a custom circuit board containing a low-cost embedded TDR frontend that can be connected to a security mesh specimen to measure its response, creating a fingerprint of the mesh. In a standard PCB manufacturing process, we construct a security mesh with a ground plane underneath that works similarly to previous work [16, 30, 12]. When viewed in the microwave domain, such meshes constitute what is essentially a delay line. Security meshes commonly use a pair of two traces to capture short circuit conditions between adjacent traces, which we treat as a differential pair for improved resiliency against electromagnetic interference. We constructed our frontend such that it excites the two traces differentially, but allows for both single-ended and differential measurements.

In an intact mesh, we expect our frontend to record no significant reflections until the stimulus pulse has traversed the mesh's traces both ways, at which point we expect a large response whose polarity and amplitude depend on the termination on the far end of the mesh. In our prototype circuit, we made this termination configurable to expand the range of possible measurement configurations and to enable self-calibration of the circuit.

Tampering with the mesh is likely to cause an impedance discontinuity. Cuts of one or both traces or a short circuit between both traces will result in a total reflection of the incident pulse at the location of the fault, which our circuit will easily detect as the delay of the response changes. However, beyond these simple cases, our approach can also detect more subtle changes. For instance, a short circuit between two points along the same mesh trace will result in a change in delay along this trace. Furthermore, even just probing a mesh trace with an oscilloscope probe will add the probe's input capacitance, resulting in an impedance step. The TDR approach is thus able to not only detect but distinguish and even localize several types of faults or attacks in a mesh.

3.2 Signal Routing

The stimulus pulse in a TDR-based design is a high-speed signal not unlike any other high-speed data or radio signal. This enables the use of signal switch and multiplexer ICs marketed for RF or high-speed data bus applications. Due to their mass-market applications, such devices are inexpensive. Using a tree-shaped topology of multiplexers, several mesh segments can be monitored by a single frontend, enabling the monitoring of arbitrarily large volumes. As a proof of concept, in our prototype we implemented software-controllable flipping of the mesh using TMUXHS4212 bus multiplexers.

3.3 Typical System Design and Threat Model

A typical system design for an HSM with TDR-based tamper sensing meshes would consist of a PCB assembly containing payload components as well as the mesh monitoring circuit. Tamper-sensing meshes made from rigid or flexible PCBs would enclose this PCB assembly from all directions. In this paper we propose meshes that have a ground plane, which would be on the outer side of the mesh PCBs and shield the system against electromagnetic interference. Mesh monitoring would be battery powered and would periodically check for tamper attempts.

We consider an attacker motivated to extract the payload's secrets. Self-destruction by deleting secrets would suffice as tamper response against this type of attacker. Such an attacker might want to probe parts of the payload circuit using either conventional electrical contacts or using electromagnetic near-field probes that must be placed right on top of the feature to be probed. An attacker might further attempt to manipulate the payload circuit, such as by removing capacitors to enable a later power side-channel attack. In preparation for an optical fault-injection attack, an attacker might attempt decapsulating some of the payload circuit's ICs either using laser ablation or using chemical etching. An attacker might also attempt fault injection attacks using either electrical contacts or electromagnetic fault injection probes near a target feature.

We consider attackers that have access to industry-standard SMD rework equipment such as microscopes, microsoldering irons, and fine tweezers. We also consider attackers that have access to more advanced equipment, such as laboratory measurement equipment like high-bandwidth oscilloscopes and waveform generators. We consider attackers with standard equipment for mechanical manipulation including precision milling machines and cutters. We do not consider bespoke attack tools, or specialized tools for large-scale industrial manufacturing such as industrial drilling machines.

4 Circuit Design and Driving Approach

A TDR can be broken down into three basic components: A source of fast stimulus pulses (or edges!), a coupler that separates stimulus pulses and their reflection at the output, and a fast ADC to capture the reflections.

Figure 2 shows a block diagram of our design[2]. At the core of our design lies an equivalent time sampling setup, where two diode bridge sampling gates alternately sample the two traces of the mesh. Since physical attacks happen on a time scale of minutes or hours, we do not need a fast acquisition rate. Equivalent time sampling uses fast

[2] Full schematics are available in this paper's supplementary material.

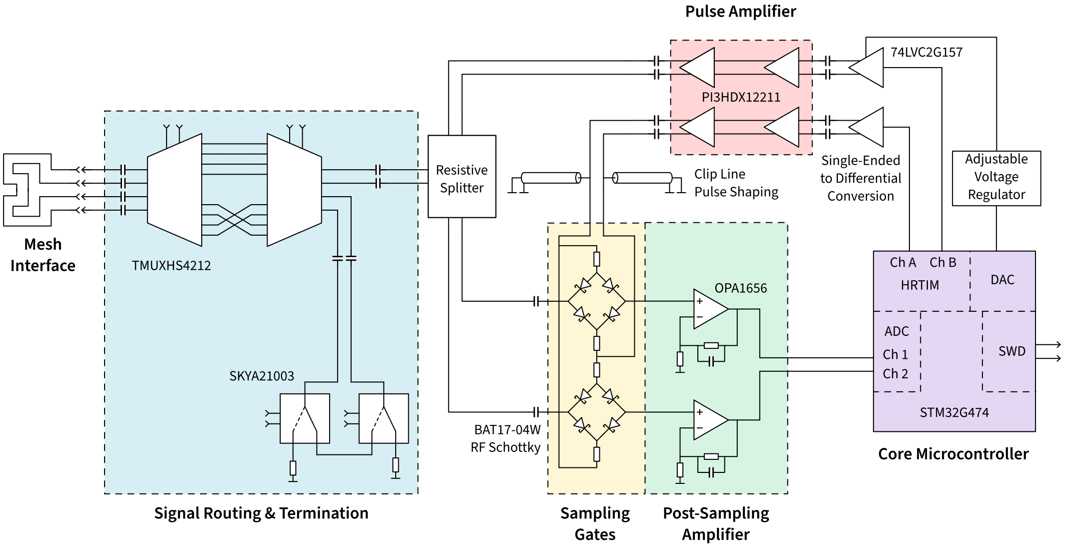

Figure 2: Block diagram of our prototype sampling TDR security mesh monitoring circuit.

sampling gates to sample a high-frequency signal at a low frequency that is suitable for direct conversion through an ADC. Using equivalent-time sampling, we can sample GHz-Scale signals at the MHz-scale sampling rate of the internal ADCs of the commodity microcontroller we use. We use two of the microcontroller's ADCs interleaved, each of which provides approximately 1.7 MSp/s at 12 bit resolution. Due to the high conversion speed of the modern ADC cores in this microcontroller, we are able to use up to 384× oversampling for increased precision.

The mesh has low insertion loss. Thanks to the resulting large amplitude of the reflection signal, the noise floor of our frontend based on commodity operational amplifiers (opamps) is below the resolution limit of the built-in ADCs of our chosen microcontroller. The main source of frontend noise stems from timing jitter between the sampling gate and the ADC due to the clock generation of the ADC, which could be reduced through firmware changes. The strong signal allows us to use a comparatively lossy but simple -6 dB resistive tee instead of a directional coupler.

We implemented the sub-nanosecond sampler using a four-diode bridge sampling gate made from commodity BAT17-04W RF Schottky diodes, which offer turn-on times better than 100 ps at 0.13 € per device at quantity 1000. In contrast to prior work [33, 14], we precisely control the timing of our ADC and avoid the need for a second sampling stage.

We base our circuit around an STM32G474RB microcontroller, 5 €-class commodity ARM microcontroller. This is a recent part, which has internal ADCs that are both higher resolution and faster than those of older parts. Furthermore, it includes a *high-resolution timer* (HRTIM) peripheral that provides better than 200 ps timing resolution through self-calibrating delay lines. We use this peripheral to produce adjustable, phase-locked stimulus and sampling pulses.

While the HRTIM peripheral provides sub-nanosecond phase adjustment, the digital outputs of the STM32G4 series are limited to a minimum transition time of $t_r = t_f = 1.7\,\mathrm{ns}$[3]. We work around this issue with two circuit tricks. First, we send the output through a fast amplifier to square up the edges to a rise time better than 500 ps. We then reduce the 10 ns minimum pulse width supported by the HRTIM peripheral by applying a clip line [42]

[3] Datasheet specification, when driving a 10 pF load [39].

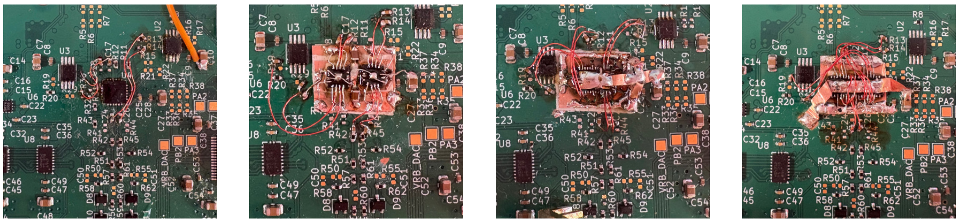

(a) 74LVC2G157 (b) MAX3748 (c) TDP0604 (d) PI3HDX12211

Figure 3: Implementation of the pulse amplifier variants of the design. Amplifiers were mounted dead bug style on copper tape and connected with 120 µm wire. Supply rails were connected with copper tape where possible to reduce impedance. MLCC power supply decoupling capacitors were placed on the copper tape to reduce loop area.

pulse forming network–i.e. we connect the amplifier's output to the load in parallel with a short, terminated transmission line stub. The length of this stub determines the pulse width.

4.1 Driver Selection

We evaluated multiple options for the pulse shaping amplifier in our design. For both sampling and stimulus, we work with fully differential signals, so Current Mode Logic (CML) devices, which are widely used in high-speed logic, are a natural fit. We settled on four parts for evaluation in this paper: A 74LVC2G157 standard logic IC, two HDMI/DisplayPort redrivers, PI3HDX12211 and TDP0604, as well as MAX3748, a limiting amplifier for optical networking. Figure 3 shows the four hand-soldered prototypes. We avoided specialty parts such as the CML-output comparators made by Analog Devices due to cost.

Standard logic ICs. As a baseline, we evaluated the 74LVC2G157 CMOS multiplexer configured to provide complementary outputs. According to manufacturer specifications, this part provides slightly faster rise and fall times than oumicrocontroller [35].

Optical Networking Chipsets. Optical transceivers use CML-output limiting amplifiers and laser drivers, some of which are still available as discrete components despite the industry moving from PCB implementations to direct bonding. We evaluated the MAX3748 limiting amplifier as a representative part from this category.

Bus Redrivers. Most modern, high-speed buses like USB 3, PCI Express, HDMI, and Display Port use CML drivers. *Redriver* ICs intended to amplify such signals to compensate for loss in connectors or cables contain amplifiers that are suitable for our application. HDMI/DisplayPort redrivers are most suitable since they can be configured as simple amplifiers, turning off any signal-dependent power saving features.

In our evaluation below, we include PI3HDX12211 and TPD0604, two inexpensive, consumer mass market redrivers[4]. Both parts have four independent channels, so only one chip is needed for the two pulse paths.

[4] PI3HDX12211 is available at 2.11 € in single quantity and less than 1.30 € at a quantity of several hundred at distributor LCSC, and TPD0604 is available at 4.72 € and 3.44 €, respectively, at distributor Mouser

Table 1: Cost breakdown of our prototype design. Prices are listed at order quantity 1000 to make prices more comparable between distributors.

Part number	Amount	Cost in €	Description
PI3HDX12211	1	1.37	Pulse amplifier
STM32G474RB	1	3.51	Main microcontroller
OPA1656	1	1.25	Sampling post-amplifier
TMUXHS4212	2	0.64	Signal routing switch
SKYA21003	2	0.49	Termination switch
74LVC2G157	2	0.15	Pulse pre-conditioning
BAT17-04W	4	0.12	Sampling gates
N/A	25	0.01	Various MLCC capacitors
N/A	25	0.01	Various resistors
		9.67	**Total**

4.2 Cost Breakdown

Table 1 shows a breakdown of the cost of the main components of our prototype, totalling less than 10 €. We did not include power supply components in this breakdown since our circuit is meant to be embedded into a payload circuit that will already have sufficient power supplies. Our design works with strong signal levels, and does not have special power supply requirements. In a practical implementation, it is unlikely that the power supply would negatively affect performance.

Due to its HRTIM peripheral, the STM32G4 microcontroller is the component of our design that is hardest to replace. However, this part can still be replaced with a wide range of FPGAs, which commonly include digitally configurable delay lines on their IO pins for signal de-skewing. For instance, the ODELAY primitive of Xilinx 7 Series FPGAs provides the same $\frac{1}{32}$ clock cycle resolution that the STM32G4 HRTIM peripheral provides while supporting higher input clock frequencies.

4.3 Measurement Principle and Scan Scheduling

The goal of a time domain reflectometer is to send a pulse into the Device Under Test (DUT)–i.e. in our application, the mesh–and to record all reflections returning from the DUT afterwards. In a security mesh with a few meters of total trace length, the time span between the pulse being sent and the last reflections arriving from the end of the mesh is in the order of tens of nanoseconds. Directly recording a response at this timescale would be infeasible in a commodity microcontroller, so we use equivalent time sampling.

As shown in Figure 2, our analog frontend contains amplifiers that produce the stimulus pulse, a sampling gate with amplifiers, and a coupler that couples the pulse into the mesh and couples the reflections back into the sampling gate. A microcontroller controls this frontend with two main signals: A stimulus pulse, and a sampling pulse. By adjusting the timing between these two pulses every time a stimulus pulse is sent, the microcontroller can sample the response at any chosen point in time. By sweeping across the whole time span, the microcontroller can reconstruct the waveform of the reflected signal at the sampling gate.

In our prototype, we sample the response once after each stimulus pulse. We conservatively decided on a sampling rate of 1 MSps across both channels of the mesh's differential pair. This sampling rate leaves some headroom to the 50 MHz Gain-Bandwidth Product (GBP) of the OPA1656 frontend opamp, as well as the 4 MSps that the ADCs can reach. The processing speed of the microcontroller allows individual control of the timing of each sampling pulse.

In our prototype, one sweep of a 141 ns time span consisting of 768 data points took 825 ms at 384× oversampling. The time span corresponds to 21 m of mesh length, which at a 200 µm pitch corresponds to a mesh area of 85 cm^2 and at a 1 mm pitch corresponds to 426 cm^2. By optimizing timing, moving oversampling processing out of the interrupt handler, and by interleaving four instead of two of the microcontroller's five ADC peripherals, the lower limit of acquisition time of a 768-point scan is 37 ms for 384× oversampling.

4.4 ADC accuracy and noise immunity

Our system uses high-frequency pulses for measurement, which inherently reject low-frequency noise components. Through our TDR approach, both the stimulus and the sampling pulses are phase-locked, functioning similarly to a lock-in amplifier. This significantly attenuates asynchronous noise. We excite the mesh with a differential signal, similar to standards such as Ethernet or HDMI. Differential signaling cancels out external interference, which tends to affect both lines equally[5].

Our front-end circuit is designed such that the analog signal entering the ADCs is strong and low in noise. Due to the high sample rate of the microcontroller's internal ADCs, we can apply extensive oversampling (384×) to enhance resolution.

5 Experimental Evaluation

We evaluated our design in two phases. In the first phase, we measured the electrical performance of our sampling circuit. The key figure in our application is the pulse generators' rise time, which determines the level of detail that we are able to extract. Since we aim at fingerprinting a connected mesh, not at performing absolute measurements, we do not need to characterize or de-embed the transfer function of our TDR frontend.

In the second phase, we evaluated the actual performance of our design on a set of 500 mesh test specimens of different layouts and structure sizes. We include detailed performance figures for a simple baseline classifier for attack detection.

5.1 Rise Time Measurement

The level of detail our frontend can extract from a mesh is limited by the rise time of the pulses it generates. We characterized this rise time both externally, using a wideband spectrum analyzer (Section 5.1.1), and through self-characterization of the circuit (Section 5.1.2). Both measurements differ because of the non-linear characteristic of the sampling Schottky pairs. Depending on the IC, our pulse generator produces output waveforms with 470 mV to 3200 mV differential voltage swing. Since the sampling diode pairs start to conduct at a combined forward voltage of approximately 300 mV, they will transition from high impedance to low impedance during a corresponding 300 mV window at the middle of the strobe pulse's edge. Thus, even if the strobe pulse shows a low-pass response with rounding at both ends, as long as its slew rate $\frac{\mathrm{d}V}{\mathrm{d}t}$ during the zero crossing is fast enough, the pulse will still result in a sharp turn-on knee of the sampling diodes.

5.1.1 Stimulus Pulse Rise Time at the Mesh

To determine the rise time of our frontend's pulse generator, we measured the stimulus output at the mesh interface using a Keysight N9020A MXA 26.5 GHz signal analyzer[5]. All measurements were taken with the prototype's mesh interface connected to the spectrum

[5] The spectrum analyzer used significantly exceeded the capabilities of the fastest oscilloscopes we had access to, so it was the more appropriate choice of measurement instrument.

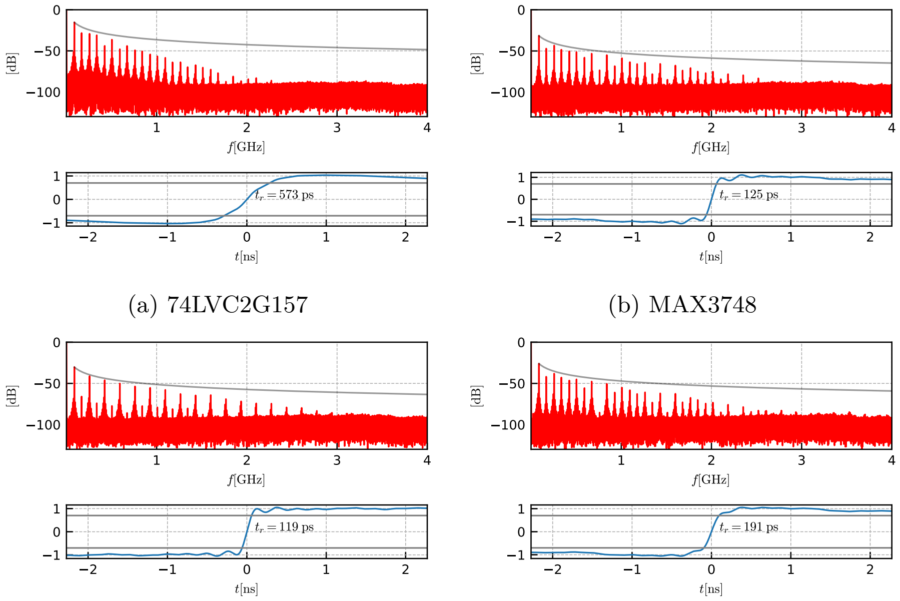

(a) 74LVC2G157

(b) MAX3748

(c) TDP0604

(d) PI3HDX12211

Figure 4: Spectrum measurements and reconstructed time domain edge shape of the stimulus pulse measured at the mesh interface for each of the four driver ICs, captured using a spectrum analyzer. Vertical scale shows arbitrary units. Spectrum plots include a $\frac{1}{f}$ reference curve indicating an ideal infinite-bandwidth square wave.

analyzer through a bias tee configured for DC blocking followed by a 20 dB attenuator for protection.

Figure 4 and Table 2 show the resulting measurements both in the frequency domain (upper traces), and projected back into the time domain (lower traces) along with measured rise times. As expected, the bare 74LVC-series logic gate has the slowest rise time at approximately 500 ps. All three amplifier variants we implemented showed significantly improved rise time, with the PI4HDX12211 achieving below 200 ps, and the other two showing around 120 ps. MAX3748 and TDP0604 only achieved a low output signal amplitude, which stems from a combination of them having low output amplitude by design and of our circuit loading their outputs heavily. Since their amplitude is only marginally within the knee region of the RF Schottky diodes used in the sampling bridges, in these variants, the sampling gates end up slower than the raw pulse rise time value alone would suggest.

5.1.2 Self-Characterization

While a fast edge is a necessary component for a fast sampling gate, the concrete speed of the sampling gate also depends on other factors such as the pulse's amplitude. Figure 5 shows the result of our self-characterization experiments, where we used the frontend to measure its own pulse shape representing its concrete sampling performance. In these experiments, we used 256× oversampling at 12 b ADC resolution. The plots show the voltage at the ADC input against time in ns. The absolute voltage levels are not relevant here - only the rise time is. Since we use some of these amplifiers–particularly the redriver ICs–well outside of their intended application, the actual voltage they develop across the nonlinear load that our sampling gate's diode bridge presents depends on implementation details of the amplifier's CML output stage. To maximize ADC resolution and minimize ringing, we tuned gain and bandwidth of each post-sampling amplifier for each IC. Ringing

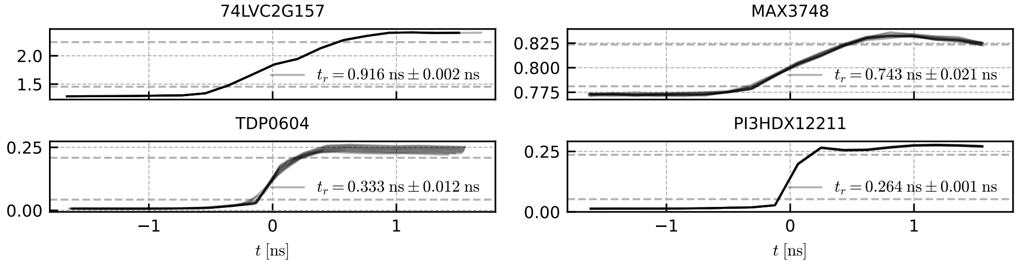

Figure 5: One edge of the stimulus pulse with no mesh connected measured by the board itself, using different amplifier ICs. For each IC, ten traces are shown. The vertical scale is in Volts at the sampling amplifier output.

Table 2: Single-ended stimulus edge rise times for different amplifier ICs. The single-ended rise times of both positive and negative half of the differential pair have been averaged. External measurements are from Figure 4, measuring the stimulus pulse at the mesh interface. V_{pp} measurements are taken at the mesh interface. Effective slew rates are calculated from the external measurements and pulse Vpp.

IC	74LVC2G157	MAX3748	TDP0604	PI3HDX12211
t_r **(Self-Characterization)**	916 ps	743 ps	333 ps	264 ps
t_r **(Stimulus at Mesh)**	573 ps	125 ps	119 ps	191 ps
Stimulus Pulse V_{pp}	1600 mV	236 mV	254 mV	430 mV
Effective Slew Rate	2.79 V ns^{-1}	1.89 V ns^{-1}	2.13 V ns^{-1}	2.25 V ns^{-1}

in the amplifier output leads to jitter in the ADC's sampling period to directly feeding through to the ADC output value. Since in STM32 MCUs, the ADC is clocked independently of the rest of the system, its sampling timing is poorly controlled and this jitter causes a significant error unless the amplifier is well-compensated.

Table 2 shows rise times calculated from each trace, averaged across both traces of the differential pair. Our results show that the optical networking limiting amplifier produces slower edges than the measurements from Figure 4 would suggest. We suspect that this is caused by its low output amplitude resulting in part from its specifications and in part from a poor match between its CML output structure and the nonlinear impedance presented by the sampling diode bridges. Surprisingly, even the 74LVC2G157 baseline unit has a rise time of less than 1 ns. We estimate that this is caused by the large output voltage swing of this part, going from ground to its V_{CC} at 3.3 V. Due to the construction of our sampling gate, its switching happens in the short period between its input differential voltage crossing zero and it rising above the combined forward voltage of the Schottky diodes. Thus, while the 74LVC might produce slow edges overall, its large output swing results in a high slew rate in the critical region around the zero crossing.

We observed the best result overall with the PI3HDX12211 redriver, resulting in a rise time of 264 ps. In this test specimen, we fed the pulse through the amplifier twice since we had two unused channels, and we used 200 ps clip lines on the amplifier's output for pulse shaping. We only used clip lines here and for TDP0604 since the other amplifiers' output did not contain sufficient harmonic content.

5.2 Mesh Specimen Characterization

To measure the practical performance of our prototype, we created a set of tamper sensing mesh test specimens. Each specimen contains four separate meshes with the same area. Table 3 shows the design specifications. Each specimen contains four separate meshes on

Table 3: Specifications of mesh test specimens used in the experiments in this paper. Approximate signal delays were calculated using wave velocity $v = \frac{c}{\sqrt{\epsilon_r}} \approx \frac{c}{2}$ [49] assuming $\epsilon_r \approx 4$ [27] for the test specimens' FR-4 substrate.

Mesh	1	2	3	4
Size	35 × 70 mm	35 × 70 mm	35 × 70 mm	35 × 70 mm
Area	$24.5\,\mathrm{cm}^2$	$24.5\,\mathrm{cm}^2$	$24.5\,\mathrm{cm}^2$	$24.5\,\mathrm{cm}^2$
Trace width	150 µm	200 µm	300 µm	500 µm
Trace spacing	150 µm	200 µm	300 µm	500 µm
Trace pitch	300 µm	400 µm	600 µm	1.00 mm
Trace length	1.07 m	1.93 m	2.86 m	3.86 m
Approximate Delay	7.1 ns	13 ns	19 ns	26 ns

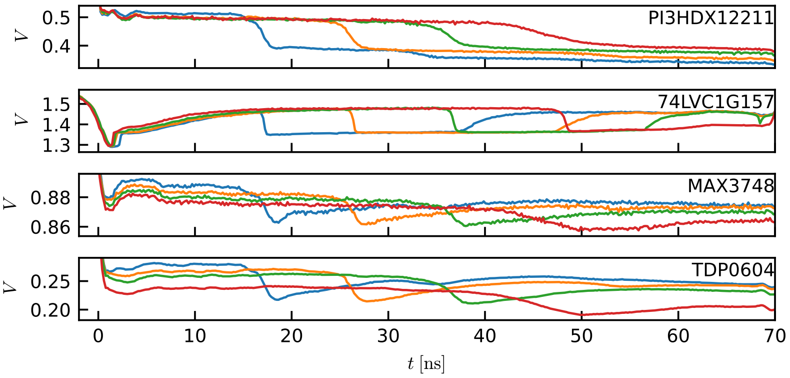

Figure 6: TDR responses captured by the microcontroller's internal ADCs with each of four candidate pulse amplifier ICs and four test meshes. The shown time range covers the primary reflection of the stimulus pulse's falling edge. For clarity, only one channel of the differential response is shown.

the outer layers of a four-layer, 1.0 mm thickness PCB, two equal-size meshes on each side. The inner layers were used as ground. Figure 6 shows the results of a baseline measurement of each mesh using each design variant. The step response resulting from an edge entering the mesh and its reflection arriving back at the start after traversing the mesh back and forth is clearly visible.

We validated the results from Figure 6 by calculating speed of light in our mesh specimen's substrate based on them. The resulting measurements are shown in Table 4. All amplifier configurations yield comparable measurements of approximately $1.6\,\mathrm{m\,s}^{-1}$, which corresponds with the expected signal propagation velocity in FR-4 PCB material of $1.5 \times 10^8\,\mathrm{m\,s}^{-1}$ [49, 27].

The graphs in Figure 6 show a dispersion effect that increasingly rounds off the trailing edge of the response with longer mesh lengths. This effect stems from higher-frequency components coupling into adjacent trace segments further up or down the mesh, spreading high-frequency components of the response signal out throughout time. This effect is less visible in the 74LVC measurements, which we suspect is a result of this variant's large pulse amplitude, which enables reflected response components to forward-bias the sampling gate's diode bridges, resulting in amplitude clipping.

From this dispersion effect follows a key point for the design of practical security meshes: To increase the temporal resolution of TDR mesh monitoring, meshes should be broken up into segments that are multiplexed through signal switching.

Table 4: Speed of light and time offset calculated from delays read from the graphs in Figure 6. c is the speed of light determined by linear fit.

	Mesh				
Pulse amplifier IC	1	2	3	4	Calculated speed of light c
PI3HDX12211	16.9 ns	26.0 ns	36.4 ns	46.1 ns	$1.59 \times 10^{8}\,\mathrm{m\,s^{-1}}$
74LVC2G157	17.1 ns	26.4 ns	36.6 ns	48.2 ns	$1.55 \times 10^{8}\,\mathrm{m\,s^{-1}}$
MAX3748	17.2 ns	26.4 ns	36.6 ns	45.6 ns	$1.59 \times 10^{8}\,\mathrm{m\,s^{-1}}$
TDP0604	17.0 ns	26.2 ns	36.5 ns	45.8 ns	$1.59 \times 10^{8}\,\mathrm{m\,s^{-1}}$

5.3 Classification performance

To evaluate the practical performance of our system, we captured approximately 1250 measurement series under a variety of environmental and attack conditions and evaluated its performance using a simple template-matching classifier. In each measurement series, we captured 7 differential traces with 2×768 points per trace. One differential trace served as a calibration reference with the multiplexers configured to disconnect the mesh. The other six traces cover each of open circuit, short circuit, and matched load termination measuring each of the two traces of the mesh once from each of both ends for 12 channels total ($\{\text{open}, \text{short}, \text{load}\} \times \{\text{forward}, \text{reverse}\} \times \{\text{mesh trace A}, \text{mesh trace B}\}$).

Our classifier is designed to compare two measurement series and produce a scalar score indicating their similarity. A simple threshold can then be applied on the similarity score to decide the class. Type 1 and type 2 error rates can be tuned by adjusting this threshold.

Our classifier proceeds in four steps: B-spline smoothing, per-channel Pearson Correlation Coefficient, averaging all channel results, and applying a threshold. B-spline smoothing serves as a low-pass filter, evening out random noise. We calculate the Pearson Correlation Coefficient for each measurement channel separately, producing a vector with 12 entries. We average the components of this vector to a single, scalar similarity score.

5.3.1 Interpreting these performance plots

Figure 7 shows the similarity score of multiple intact meshes. For each performance measurement, we show the similarity scores for each pair of measurements as a matrix, with each measurement appearing once in each row and column. High values indicate similarity, low values indicate differences. We show the baseline measurement set in the top left quadrant of the plot (1), and the experiment set bottom right (4), separated by white lines. Uniform color within the top left quadrant (1) indicates high similarity between baseline measurements. Nonuniform color in the bottom right (4) is expected, and indicates that mutliple experiment (attack) measurements are unlike each other. Classification performance is indicated by the top right (2) and bottom left (3) quadrants, which indicate misclassification probability. Misclassification is likely when the top left (1) and top right (2) quadrants look alike. Misclassification is less likely the more they differ.

Under each figure, we give the False Negative Rate (FNR) when the threshold is adjusted for a False Positive Rate (FPR) of 0.1% as a reference point[6]. We also provide the Crossover Error Rate (CER) at which for some threshold FPR is equal to FNR. We calculate all error rates assuming the similarity scores are normally distributed. We chose a reference point of 0.1% FPR since it allows for a meaningful comparison based on the hundreds of measurements our data is based on. In a practical application, the end-to-end FPR of the alarm system would need to be significantly lower, probably in the range from 10^{-12} to 10^{-9} for a Mean Time Between Failures (MTBF) of several years. A

[6] We denote the rate of missed alarms as FNR and the false alarm rate as FPR.

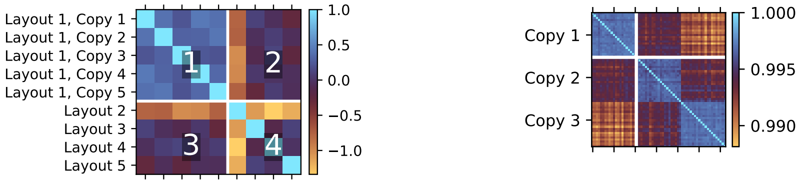

(a) Five copies of the same layout compared to four other layouts. FNR 18% at 0.1% FPR, CER=8.3%.

(b) Three identical copies, 20 measurements each. FNR 1.7% at 0.1% FPR, CER=1.1%.

Figure 7: Similarity matrices of measurement series on intact meshes.

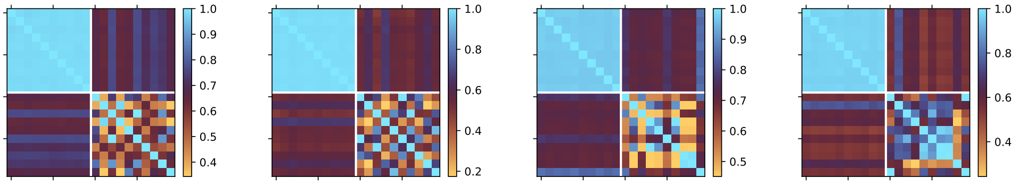

(a) One trace interrupted, p=0.3 mm. FNR 0.0% at 0.1% FPR, CER=0.0%.

(b) Both traces shorted, p=0.3 mm. FNR 0.0% at 0.1% FPR, CER=0.0%.

(c) One trace interrupted, p=0.4 mm. FNR 0.0% at 0.1% FPR, CER=0.0%.

(d) Both traces shorted, p=0.4 mm. FNR 0.0% at 0.1% FPR, CER=0%.

Figure 8: Similarity matrix of 10 intact and 10 modified meshes with two pitch sizes under two different attack scenarios: An interrupted trace, and both mesh traces shorted.

practical system would likely include additional components filtering the output of our proposed baseline classifier analyzing not just the last, but multiple previous measurements. Experimentally evaluating a classifier to this degree of precision would require a large-scale experiment to account for the long tail of the error distribution.

Figure 7a compares several copies of the same mesh (top left quadrant, 1) to four variants that have the same pitch and area, but different randomized layout of the traces (bottom right). Our classifier can distinguish mesh layouts with a 18% FNR at 0.1% FPR.

The variance between samples of the baseline group in Figure 7a alerted us to the possibility that while all mesh samples of the same layout were supposed to be identical copies, our measurement circuit might be sensitive enough to pick up on manufacturing variations from one copy to another in a PUF-like manner. To evaluate this scenario, in Figure 7b we show the result of repeated measurements of three copies of the same mesh. The measurements were taken interleaved $(1, 2, 3, 1, 2, \ldots)$ to exclude systematic errors. We found our system can indeed distinguish multiple copies of the same mesh at a 1.7% FNR at 0.1% FPR. We leave a detailed analysis of this effect to future work. For the scope of this paper, the presence of this effect indicates good performance of our design, and increases the detection efficiency of our approach.

5.3.2 Basic attacks

Figure 8 shows the performance of our classifier under the two basic attack scenarios of an interrupted trace, and a short circuit between the mesh's differential traces. Such attacks lead to large changes in the location of the reflected pulse edge, resulting in 0% Crossover

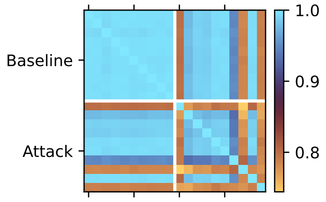

Figure 9: Similarity matrix of several mesh specimens that have one trace shorted to an adjacent location on the same trace. Classification FNR 23% at 0.1% FPR, CER=22%.

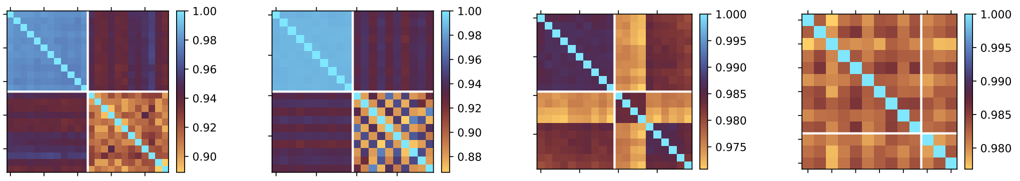

(a) Oscilloscope probe contacting mesh. FNR 0.0% at 0.1% FPR, CER=0.0%.

(b) Soldering iron touching mesh. FNR 0.0% at 0.1% FPR, CER=0.0%.

(c) 30mm wire soldered to mesh. FNR 9.6% at 0.1% FPR, CER=6.7%.

(d) Baseline vs. experiment specimens with no attack.

Figure 10: Classifier performance under advanced attack scenarios.

Error Rate.

5.3.3 Trace shortening

Figure 9 shows classification results when one trace is short circuited to another location within the same trace. Here, the resulting distortion in response shape is harder to detect. Depending on the length of the shorted-out section, the timing skew such modifications introduce may be as little as a few picoseconds. For some samples which have longer sections of mesh trace shorted out, this attack is easy to distinguish, but for others, our classifier cannot distinguish it leading to an overall FNR of 18% at 0.1% FPR, with some specimens reliably detected, and others never detected.

5.3.4 Advanced attacks

Figure 10 shows our classifier's performance under conditions similar to actions an attacker would perform during an attack: An oscilloscope probe[7] touching one mesh trace (Figure 10a), a soldering iron touching one mesh trace (Figure 10b), and a mesh where one trace has a $l = 30\,\mathrm{mm}, d = 120\,\mu\mathrm{m}$ piece of copper wire soldered to one trace (Figure 10a). Our classifier is able to clearly distinguish the probing and soldering iron cases at 0% FNR, with a maximum of 9.6% FPR at 0.1% FNR in the soldered wire case.

5.3.5 Patching attacks

PCB tamper sensing meshes are susceptible to industry-standard PCB rework techniques. If we assume a standard PCB process with 100 µm trace/space design rules, a drilling attack targeting a 300 µm hole size requires cutting and patching at least one trace [17].

[7] Part number Rigol PVP3150.

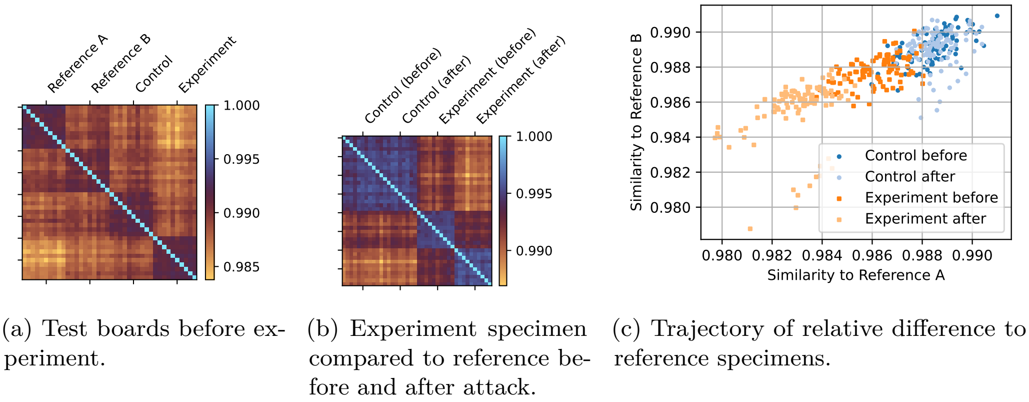

(a) Test boards before experiment.

(b) Experiment specimen compared to reference before and after attack.

(c) Trajectory of relative difference to reference specimens.

Figure 11: Classifier performance under a patching attack that bridges a short gap within a mesh trace using wire.

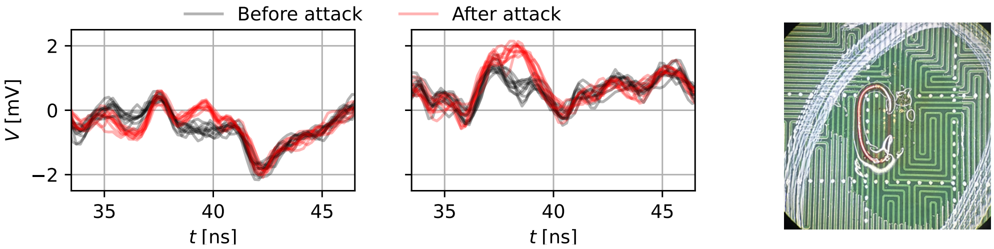

Figure 12: The mesh response under a manipulation attack patching across a drill location for a 300 µm drill, as captured by the microcontroller's ADCs. The mesh pitch is 300 µm. B-spline smoothing was applied for readability.

We performed such an attack on a set of 300 µm pitch meshes. Figure 12 shows our modification and the resulting change in the time-domain response.

Figure 11 shows the classification result of this attack. To extract the subtle effect of this attack, we measured two reference specimens, one control, and one experiment specimen twice: Once before the attack, and once after. Measurements were interleaved and repeated 10 times. Factors such as temperature drift can be excluded by comparing both control and experiment measurements against the two references before and after the modification. Figure 11a shows the four samples before the attack, exhibiting the same subtle PUF-like effect that we described in Section 5.3. Since we peform both before and after measurements on the same sample, we can separate this effect from the effect of the attack. Figure 11b compares both control and experiment samples before and after the attack, and shows a clear change in the experiment sample during the attack. Figure 11c plots the similarity scores of both samples to each of the two reference samples. We can see that the control distribution stays in one place, while the experiment distribution shifts.

Based on the above results, we peformed a larger-scale experiment using ten interleaved measurements each of seven samples with patches applied compared against baseline measurements taken before and after measuring the experiment samples.. Figure 13 shows the results of this experiment, resulting in a FNR of 71.5% at 0.1% FPR. Since such patches only affect few data points along the reflection response, we included a variant of our classifier that uses the maximum difference across all channels instead of the averaged Pearson Correlation Coefficient to improve sensitivity to the subtle, localized effects of

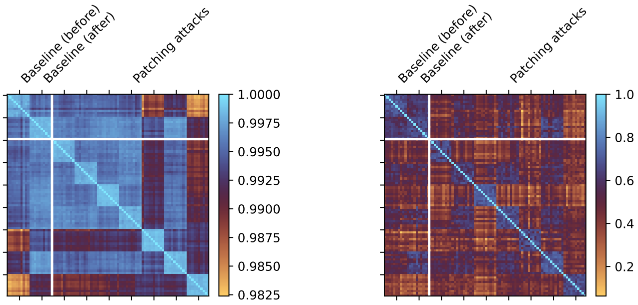

(a) Micro-soldering patching attack. FNR 71.5% at 0.1% FPR, CER=29%.

(b) *maximum* classifier variant. FNR 51.1% at 0.1% FPR, CER=15%.

Figure 13: Classification performance in a larger-scale experiment using 10 measurements each of 7 samples with traces patched through micro-soldering.

such patches. Using this classifier variant, FNR improves to 51.1%, detecting half of all attack attempts in a single measurement when fixing the false alarm rate at 0.1%. In a practical application, detection rates would be higher since the system would be able to observe the entire process of patching. As shown in Section 5.3.4, soldering for instance is highly detectable, while here we only benchmark a momentary snapshot after the patch was completed.

5.3.6 Environmental susceptibility

Figure 14 shows the results of a series of experiments evaluating the effect of environmental factors such as handling or electromagnetic interference on our measurements. Figure 14a shows our measurements exhibit little time drift (CER=60%). Figure 14b shows that touching the mesh is easily detected (FNR=0%), but the system is insensitive to touching other parts of the circuit. We classify touching the mesh as an attack since the mesh would be shielded from touch by the ground plane in a practical scenario (cf. Section 3.3).

As shown in Figure 14c, heating the mesh distors its measurements (FNR=0.6%, CER=0%). Figure 15 shows the difference caused by heating the mesh to 70 °C in the time domain. This temperature dependence stems from the resistance of the mesh's copper traces increasing with temperature, and the dielectric properties of the FR-4 PCB substrate changing. Both dielectric constant and dissipation factor of FR-4 change with temperature [37, 13]. The increase in copper resistance causes a shift of the response curve. An increase in the dielectric dissipation factor affects the slope of the difference in Figure 15 since pulse energy is dissipated more the longer the pulse travels through the material. A change in dielectric constant moves the response's trailing edge in time, with the pulse propagating slightly slower at high temperature.

Since these effects are consistent with physical predictions and only reach problematic levels at large temperature differences, it would be possible to design a classifier that is insensitive to temperature effects. Furthermore, given the predictable, physical nature of these effects, they could also be compensated before classification in the digital domain based on a temperature measurement.

Besides temperature, other environmental factors such as electromagnetic interference could theoretically also influence our measurements. Although our system's equivalent-time sampling setup inherently cancels out EMI since it is not synchronous to the sampling clock, the setup is unshielded so we verified its actual susceptibility in several scenarios. Figure 16 shows the result of these measurement series. For comparison, we included

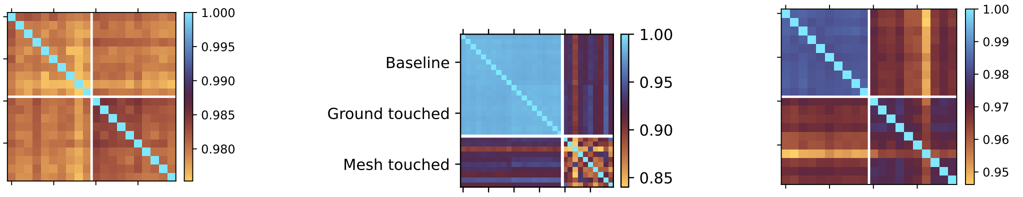

(a) Time drift (2.5h). FNR 100% at 0.1% FPR, CER=61%.

(b) Touch sensitivity. FNR 0.0% at 0.1% FPR, CER=0.0%.

(c) Mesh heated (70 °C).

Figure 14: Classification results of the same mesh under various environmental factors.

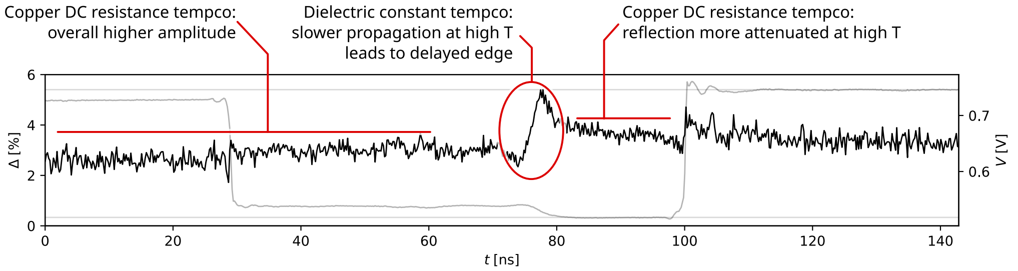

Figure 15: The effect of heating on a time-domain trace. One of 12 channels shown. Gray: Raw data. Black: Relative difference between hot and cool cases.

several measurements from Figure 13. From these figures, we can see that there are some environmental effects, but these effects are small even when compared against a subtle attack like a patching attack with the classification performance remaining approximately constant at 69.0% FNR at 0.1% FPR and a slightly reduced CER of 20%.

5.4 Countermeasures

As shown above, PCB security meshes can be manipulated through micro-soldering. Keeping the modifications as physically small as possible, their impact on TDR response can potentially be kept below detection thresholds of our single-shot baseline classifier. However, even with such a simple classifier, the entire attack would have to be carried out without raising an alarm, e.g. by touching the mesh or contacting a trace with the soldering iron. Soldering would have to be done using a minimal amount of solder as well as a bespoke, insulated soldering iron tip. While manufacturing such a tool out of a material like sintered ceramic is conceivable, to our knowledge, no such tool exists on the market.

Furthermore, the actual drilling would have to happen with a dielectric drill bit, placing special attention on evacuating conductive copper chips before they can create short circuits to nearby traces. Again, it is conceivable that such a tool could be manufactured, but to our knowledge, such a tool is not currently available as a standard component on the market.

Finally, any probes penetrating the mesh would have to be placed such that their presence in the vicinity of the mesh traces does not disturb the TDR response. Modifications would have to be carried out with great care, likely using micromanipulators or similar specialized equipment.

The PCI PTS HSM DTR standard [32] contains a useful framework for thinking about

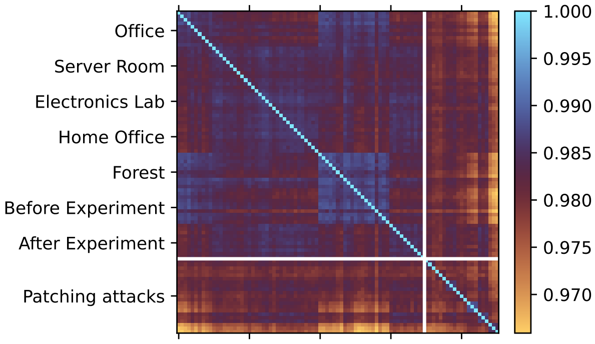

Figure 16: Classifier similarity scores of measurements in different environments, 10 measurements each. For scale, measurements from Figure 13 are included on the bottom/right. FNR 69.0% at 0.1% FPR, CER=22%.

attacker capabilities. Applying their taxonomy, our monitoring system raises the skill level required for a patching attack from a *skilled* attacker to an *expert* attacker, and the equipment requirement from *standard* equipment to *bespoke* equipment.

6 Future Work

Advanced attack classification. While we proposed a simple baseline classifier, there is a large parameter space for more advanced designs. For instance, a classifier could apply machine learning techniques to adapt to the response of a particular mesh, learn its benigh behavior under temperature changes, and dynamically schedule sample timing to focus attention on the parts of the response signal that are most susceptible to attacks. Moving from a single-shot classifier that only observes measurements in isolation to a more advanced approach that considers the full history of measurements during the mesh's lifetime would also likely improve performance.

Auxiliary applications. The low-cost, embedded TDR frontend presented in this paper could be used for other monitoring tasks from tamper sensing to system health monitoring. For instance, Vai et al. [44] propose checking the integrity of a PCBA using an external Vector Network Analyzer (VNA) attached to test points on the PCBA's Power Distribution Network (PDN). TDR can produce fingerprints similar to a VNA and it would be interesting to measure parts of the secure subsystem other than its security mesh using our TDR frontend.

Characterization of PUF-like effects. In Section 5.3, we have described a PUF-like effect, where our classifier was able to distinguish supposedly identical copies of the same mesh. It would be interesting to precisely characterize this effect and its dependence on factors such as the chosen PCB manufacturer, and to quantify if it indeed rises to the level of a PUF in entropy and repeatability.

7 Conclusion

In this paper, we presented a design for a low-cost frontend for integrity monitoring of security meshes in applications such as HSMs based on the principles of sub-nanosecond Time Domain Reflectometry. Our design repurposes an inexpensive HDMI redriver IC and uses a microwave clip line to form fast pulses for TDR sampling. Our design creates a detailed fingerprint of the intact mesh's condition that not only captures the length of the mesh's traces but that can distinguish copies of the same mesh.

We have demonstrated our prototype circuit's capability to reliably detect and distinguish a wide range of practical attacks with no classification erros in most attack classes, and a worst-case FNR of 71.5% at 0.1% FPR when detecting tiny, micro-soldered patch wires.

Compared to the state of the art, our approach enables the monitoring of larger meshes, at higher sensitivity and lower cost. Our is easy to replicate, does not require any specialized or custom components, and unlocks high-security applications for security meshes made using low-cost, standard PCB manufacturing processes.

Availability

This is version `v3-final-1-gb4dc582` of this paper, generated on October 20, 2025. The git repository with the LaTeX source for this paper, all hardware design files, and firmware and analysis source code can be found at:

https://git.jaseg.de/sampling-mesh-monitor.git

References

[1] (US) National Institute of Standards and Technology. *Security Requirements for Cryptographic Modules*. Tech. rep. Federal Information Processing Standard (FIPS) 140-2. U.S. Department of Commerce, Dec. 2002. DOI: 10.6028/NIST.FIPS.140-2.

[2] R. Anderson et al. “Cryptographic Processors-A Survey”. In: *Proceedings of the IEEE* 94.2 (Feb. 2006), pp. 357–369. ISSN: 1558-2256. DOI: 10.1109/JPROC.2005.862423.

[3] Ross Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 1st ed. Wiley, Dec. 2020. ISBN: 978-1-119-64278-7. DOI: 10.1002/9781119644682.

[4] G. Bencivenni et al. “A Time Domain Reflectometer with 100 Ps Precision Implemented in a Cost-Effective FPGA for the Test of the KLOE-2 Inner Tracker Readout Anodes”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 698 (Jan. 2013), pp. 185–191. ISSN: 0168-9002. DOI: 10.1016/j.nima.2012.10.023.

[5] Eric Bogatin. *Signal and Power Integrity, Simplified*. Third edition. Boston: Prentice Hall, 2018. ISBN: 978-0-13-451341-6.

[6] William L. Brodsky et al. “Tamper-Respondent Assembly With Flexible Tamper-Detect Sensor(s) Overlying In-Situ-Formed Tamper-Detect Sensor”. 10,327,329 B2. June 2019.

[7] Britton Chance et al., eds. *Waveforms*. Vol. 19. MIT Radiation Laboratory. New York: McGraw-Hill, 1949.

[8] Raymond O. Chock and Mark Hess. “Point of Sale Terminal Having Pulsed Current Tamper Control Sensing”. US7551098B1. June 2009.

[9] François Dupont et al. “A Miniaturized and Ultra-Low-Power Tamper Detection Sensor for Portable Applications”. In: *IEEE Sensors Journal* 22.5 (Mar. 2022), pp. 4524–4533. ISSN: 1558-1748. DOI: 10.1109/JSEN.2022.3143656.

[10] Daisuke Fujimoto et al. “A Demonstration of a HT-Detection Method Based on Impedance Measurements of the Wiring Around ICs”. In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 65.10 (Oct. 2018), pp. 1320–1324. ISSN: 1549-7747, 1558-3791. DOI: 10.1109/TCSII.2018.2858798.

[11] Kathrin Garb et al. “FORTRESS: FORtified Tamper-Resistant Envelope with Embedded Security Sensor”. In: *2021 18th International Conference on Privacy, Security and Trust (PST)*. Dec. 2021, pp. 1–12. DOI: 10.1109/PST52912.2021.9647783.

[12] Kathrin A Garb. “Tamper-Sensitive Design of PUF-Based Security Enclosures”. PhD thesis.

[13] Scott Hinaga et al. “Thermal Effects on PCB Laminate Material Dielectric Constant and Dissipation Factor”. In: *IPC Apex Expo*. 2010.

[14] Hubert Houtman. “1-GHz Sampling Oscilloscope Front End Is Easily Modified”. In: *Electronic Design* 48.19 (Sept. 2000), pp. 175–176. ISSN: 0013-4872.

[15] *HP 187B Dual-Trace Vertical Amplifier Operating and Service Manual*. Hewlett-Packard Company, 1962.

[16] Vincent Immler et al. “B-TREPID: Batteryless Tamper-Resistant Envelope with a PUF and Integrity Detection”. In: *2018 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. Apr. 2018, pp. 49–56. DOI: 10.1109/HST.2018.8383890.

[17] Vincent Immler et al. “Secure Physical Enclosures from Covers with Tamper-Resistance”. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (Nov. 2018), pp. 51–96. ISSN: 2569-2925. DOI: 10.46586/tches.v2019.i1.51-96.

[18] *ISO/IEC 24759:2025*. https://www.iso.org/standard/82424.html.

[19] M. Kahrs. “50 Years of RF and Microwave Sampling”. In: *IEEE Transactions on Microwave Theory and Techniques* 51.6 (June 2003), pp. 1787–1805. ISSN: 1557-9670. DOI: 10.1109/TMTT.2002.806934.

[20] Donghwan Lee, Jinho Sung, and Jaehong Park. “A 16ps-Resolution Random Equivalent Sampling Circuit for TDR Utilizing a Vernier Time Delay Generation”. In: *2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515)*. Vol. 2. Oct. 2003, 1219–1223 Vol.2. DOI: 10.1109/NSSMIC.2003.1351912.

[21] Alan Henry Leek and Jace Hunter Hall. “Tamper Detection”. 10,925,154 B2. Feb. 2021.

[22] Yibiao Lu et al. “Correlated Randomness Teleportation via Semi-trusted Hardware—Enabling Silent Multi-party Computation”. In: *Computer Security – ESORICS 2021*. Ed. by Elisa Bertino, Haya Shulman, and Michael Waidner. Vol. 12973. Cham: Springer International Publishing, 2021, pp. 699–720. ISBN: 978-3-030-88427-7 978-3-030-88428-4. DOI: 10.1007/978-3-030-88428-4_34.

[23] Saleh Khalaj Monfared, Tahoura Mosavirik, and Shahin Tajik. “LeakyOhm: Secret Bits Extraction Using Impedance Analysis”. In: *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’23. New York, NY, USA: Association for Computing Machinery, Nov. 2023, pp. 1675–1689. ISBN: 979-8-4007-0050-7. DOI: 10.1145/3576915.3623092.

[24] Tahoura Mosavirik, Patrick Schaumont, and Shahin Tajik. “ImpedanceVerif: On-Chip Impedance Sensing for System-Level Tampering Detection”. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (Nov. 2022), pp. 301–325. ISSN: 2569-2925. DOI: 10.46586/tches.v2023.i1.301-325.

[25] Tahoura Mosavirik and Shahin Tajik. “BackMon: IC Backside Tamper Detection Using On-Chip Impedance Monitoring”. In: CCS ’24: ACM SIGSAC Conference on Computer and Communications Security. Salt Lake City UT USA: ACM, Nov. 19, 2024, pp. 68–77. ISBN: 979-8-4007-1235-7. DOI: 10.1145/3689939.3695784.

[26] Tahoura Mosavirik et al. “Silicon Echoes: Non-Invasive Trojan and Tamper Detection Using Frequency-Selective Impedance Analysis”. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* 2023.4 (Aug. 2023), pp. 238–261. ISSN: 2569-2925. DOI: 10.46586/tches.v2023.i4.238-261.

[27] Stephen J. Mumby and Jih Yuan. “Dielectric Properties of FR-4 Laminates as a Function of Thickness and the Electrical Frequency of the Measurement”. In: *Journal of Electronic Materials* 18.2 (Mar. 1989), pp. 287–292. ISSN: 0361-5235, 1543-186X. DOI: 10.1007/BF02657420.

[28] Catalin Negrea and Marius Rangu. “Sequential Sampling Time Domain Reflectometer”. In: *2009 15th International Symposium for Design and Technology of Electronics Packages (SIITME)*. Sept. 2009, pp. 367–371. DOI: 10.1109/SIITME.2009.5407341.

[29] John Norton. “Tamper Detecting Cases”. US10489614B2. Nov. 2019.

[30] Johannes Obermaier et al. “A Measurement System for Capacitive PUF-based Security Enclosures”. In: *DAC ’18: The 55th Annual Design Automation Conference 2018*. San Francisco California: ACM, June 2018, pp. 1–6. ISBN: 978-1-4503-5700-5. DOI: 10.1145/3195970.3195976.

[31] Barrett B. Parsons and Jerry L. Wells. “Tamper and Radiation Resistant Instrumentation for Safeguarding Special Nuclear Material”. In: *IEEE Transactions on Nuclear Science* 24.1 (Feb. 1977), pp. 616–620. ISSN: 1558-1578. DOI: 10.1109/TNS.1977.4328751.

[32] PCI Security Standards Council. *Payment Card Industry PIN Transaction Security Hardware Security Module Modular Derived Test Requirements*. Dec. 2021.

[33] Petr Polášek. “Reflektometr v Časové Oblasti”. MA thesis. Jan. 2020.

[34] Mani Razaghi and Jesse Hill. “Tamper Detection System”. US10595400B1. Mar. 2020.

[35] Renesas Electronics Corporation. *Application Note AN-224: ALVC/LVC Logic Characteristics and Applications*. 2019.

[36] Maryam Saadat Safa and Shahin Tajik. “Near-Field Microwave Sensing for Chip-Level Tamper Detection”. In: *Sensors* 25.13 (July 5, 2025), p. 4188. ISSN: 1424-8220. DOI: 10.3390/s25134188.

[37] Pankaj Sagar and Kashif Akber. “Studies on Temperature Dependent Dielectric Properties of Some Insulators down to Liquid Helium Temperatures”. In: *Cryogenics* 141 (July 2024), p. 103865. ISSN: 0011-2275. DOI: 10.1016/j.cryogenics.2024.103865.

[38] Munehiko Sato, Ivan Poupyrev, and Chris Harrison. “Touché: Enhancing Touch Interaction on Humans, Screens, Liquids, and Everyday Objects”. In: *CHI ’12: CHI Conference on Human Factors in Computing Systems*. Austin Texas USA: ACM, May 2012, pp. 483–492. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2207743.

[39] ST Microelectronics. *STM32G474xB/C/E Datasheet*. Nov. 2021.

[40] Paul Staat et al. “Anti-Tamper Radio: System-Level Tamper Detection for Computing Systems”. In: *2022 IEEE Symposium on Security and Privacy (SP)*. May 2022, pp. 1722–1736. DOI: 10.1109/SP46214.2022.9833631.

[41] Mark Tehranipoor et al. *Hardware Security Primitives*. Cham: Springer International Publishing, 2023. ISBN: 978-3-031-19184-8. DOI: 10.1007/978-3-031-19185-5.

[42] Tektronix Inc. *Tektronix S-6 Sampling Head Instruction Manual*. Sept. 1982.

[43] Dennis Trebbels et al. “Miniaturized FPGA-Based High-Resolution Time-Domain Reflectometer”. In: *IEEE Transactions on Instrumentation and Measurement* 62.7 (July 2013), pp. 2101–2113. ISSN: 1557-9662. DOI: 10.1109/TIM.2013.2245190.

[44] Michael Vai et al. “Secure Architecture for Embedded Systems”. In: *2015 IEEE High Performance Extreme Computing Conference (HPEC)*. Sept. 2015, pp. 1–5. DOI: 10.1109/HPEC.2015.7322461.

[45] D. C. Vasile and P. M. Svasta. “Temperature Sensitive Active Tamper Detection Circuit”. In: *2017 IEEE 23rd International Symposium for Design and Technology in Electronic Packaging (SIITME)*. Oct. 2017, pp. 175–178. DOI: 10.1109/SIITME.2017.8259885.

[46] Daniel-Ciprian Vasile and Paul Svasta. “Protecting the Secrets: Advanced Technique for Active Tamper Detection Systems”. In: *2019 IEEE 25th International Symposium for Design and Technology in Electronic Packaging (SIITME)*. Oct. 2019, pp. 212–215. DOI: 10.1109/SIITME47687.2019.8990877.

[47] Daniel-Ciprian Vasile et al. “Active Tamper Detection Circuit Based on the Analysis of Pulse Response in Conductive Mesh”. In: *2017 40th International Spring Seminar on Electronics Technology (ISSE)*. May 2017, pp. 1–6. DOI: 10.1109/ISSE.2017.8000987.

[48] Erling Wesselhoff. “Tamper Responsive Sensor”. US10678957B2. June 2020.

[49] H.A. Wheeler. “Transmission-Line Properties of Parallel Strips Separated by a Dielectric Sheet”. In: *IEEE Transactions on Microwave Theory and Techniques* 13.2 (Mar. 1965), pp. 172–185. ISSN: 1557-9670. DOI: 10.1109/TMTT.1965.1125962.

[50] Huifeng Zhu et al. “PDNPulse: Sensing PCB Anomaly With the Intrinsic Power Delivery Network”. In: *IEEE Transactions on Information Forensics and Security* 18 (2023), pp. 3590–3605. ISSN: 1556-6021. DOI: 10.1109/TIFS.2023.3285490.