Bootstrapping Physical Security with Inertial Hardware Security Modules

Draft build, git revision `proof-3-13-g82ac019`

Dissertation von Jan Sebastian Götte
zur Erlangung des Grades Doktor-Ingenieur (Dr.-Ing.)
am Fachbereich Elektrotechnik und Informationstechnik
der Technischen Universität Darmstadt
Erstgutachter: Prof. Dr. Björn Scheuermann
Zweitgutachter: Prof. Dr. Shahin Tajik

Darmstadt 2026

Bibliographische Angaben:
GÖTTE, Jan Sebastian: **Bootstrapping Physical Security with Inertial Hardware Security Modules**
Darmstadt, Technische Universität Darmstadt, 2025
URN: TBD FIXME
Tag der mündlichen Prüfung: TBD FIXME

Kurzzusammenfassung

This section is a translated copy of the English abstract below.

Im Laufe der letzten Jahrzehnte habe Fortschritte in der Kryptographie sowie Techniken wie formale Verifikation den Stand der Softwaresicherheit stetig verbessert. Gleichzeitig hat das Gebiet der Hardwaresicherheit mit diesen Entwicklungen nicht Schritt halten können. Trotz Fortschritten in Teilgebieten wie der Resilienz gegenüber Seitenkanalangriffen und Physical Unclonable Functions (PUFs) ist der Stand der Technik in der Hardwaresicherheit nach wie vor auf die Verwendung mikroelektronischer Strukturen fokussiert. Solche erreichen einen Grad der Security by Obscurity, liefern jedoch keine fundierteren Sicherheitsgarantien. Systemweite Manipulationsschutzmaßnahmen werden nur vereinzelt in Geräten wie z.B. Hardware-Sicherheitsmodulen (HSMs) und Kartenzahlungsterminals eingesetzt. Insbesondere HSMs werden aufgrund ihrer hohen Kosten und geringen Rechenleistung nur in Nischenanwendungen wie z.B. der Zertifikatsausstellung im Transport Layer Security (TLS)-System sowie der Zahlungsdatenverarbeitung eingesetzt.

In dieser Dissertation führt Jan Sebastian Götte das Inertiale Hardware-Sicherheitsmodul (IHSM) ein. Das IHSM ist eine neue Architektur für Hardware-Sicherheitsmodule, die einen hoch sicheren, aktiven Manipulationsschutz bereitstellt. Gleichzeitig können mithilfe der IHSM-Technologie kryptographische Rechnersysteme von wesentlich größeren Abmessungen, Gewicht und elektrischer Leistungsaufnahme geschützt werden, als das in konventionellen HSMs möglich ist. IHSMs ersetzen die kostenintensiven und in der Herstellung aufwendigen Meshes (Manipulationserkennungsmembranen) konventioneller HSMs durch eine Konstruktion, in der Meshes aus einfachen Platinen aufgebaut werden, die bei einer hohen Geschwindigkeit um das geschützte Rechnersystem rotieren. Die Rotation dieser Meshes verhindert eine unerkannte Manipulation. IHSMs erreichen so mithilfe wensentlich einfacherer und kostengünstiger Konstruktionstechniken ein Sicherheitsniveau, das dem konventioneller Manipulationsschutzmembranen gleicht, die in spe-

zialisierten Herstellungsprozessen gefertigt werden. In der Dissertation werden die Ergebnisse einer Übersichtsstudie vorgestellt, die etwa 30 echte Implementierungen socher Meshes untersucht. In der Studie werden Kriterien für die Entwicklung sicherer Meshes abgeleitet, anhand derer das IHSM-Konzept kontextualisiert wird. Um die Notwendigkeit sicherer Hardware zu erörtern, wird in dieser Dissertation darüber hinaus eine Analyse einiger problematischer Aspekte des Hardwaresicherheitskonzeptes der Deutschen elektronischen Patientenakte vorgestellt.

Um den Weg für zukünftige, praktische Implementierungen der IHSM-Technologie zu bereiten stellt Jan Sebastian Götte Lösungen für wichtige Schlüsselprobleme der Konstruktion von IHSMs vor. Diese Lösungen umfassen ein neues Konzept für rotationssymmetrische Planarspulen für die drahtlose Energieübertragung an rotierende Empfänger, sowie ein hochpräzises und dennoch kostengünstiges Überwachungssystem für Meshes. Dieses Überwachungssystem beruht auf dem Prinzip der Zeitbereichsreflektometrie und erkennt selbst fortgeschrittene Angriffstechniken zuverlässig. In praktischen Versuchen zeigte sich, dass das System ausreichend empfindlich ist, um mehrere identische Kopien desselben Meshes voneinander zu unterscheiden, was auf PUF-ähnliche Eigenschaften hindeutet.

In der Dissertation werden zwei konkrete Anwendungsszenarien erläutert, die erst durch das größere Volumen und die höhere Leistungsaufnahme möglich werden, die die IHSM-Technologie ermöglicht. Im ersten Anwendungsszenario wird eine IHSM-geschützte Zwischenstation vorgeschlagen, um die durch physikalische Grundgesetze sonst stark eingeschränkte erreichbare Entfernung eines Quantenschlüsselaustausch (QKD)-Systems zu vergrößern. Im Rahmen dieses Anwendungsszenarios wird ein sekundäres Mesh vorgestellt, dass die Achsdurchführung des primären IHSM-Meshes zusätzlich schützt. Weiterhin wird in der Fallstudie der Entwurf eines mechanischen Trägers für diese zusätzlich geschützte Achsdurchführung vorgestellt, der das QKD-System im inneren des IHSM mit der Außenwelt über verlustarme Glasfaserleitungen verbindet.

In der zweiten Fallstudie wird ein Konzept vorgestellt, das mithilfe IHSM-geschützter, leistungsstarker Serverhardware kolokierte Secure Multiparty Computation (MPC)-Berechnungen ermöglicht. Hierzu wird IHSM-Technologie an die Anforderungen leistungsstaker Serverhardware in Größe, Leistungsaufnahme, und ableitbarer Verlustleistung angepasst. Wird MPC praktisch eingesetzt, werden Knoten über mehrere Rechenzentren verteilt

um einen Single Point of Failure zu vermeiden. Diese Verteilung führt jedoch zu geringer Netzwerkbandbreite und hohen Latenzen zwischen den MPC-Knoten, was die erreichbare MPC-Rechenleistung stark einschränkt. Durch den Einsatz von IHSMs können physisch gesicherte MPC-Knoten innerhalb desselben Rechenzentrums betrieben werden, was durch die damit erreichbare höheren Bandbreiten und geringeren Latenzen einen Leistungsbereich der MPC-Berechnungen erschließt.

Abstract

In the past decades, cryptographic advancements and techniques like formal verification have steadily improved software security. Meanwhile, the field of hardware security has not kept pace. Research has made progress in subfields such as resilience to Side-Channel Attacks (SCA) and Physical Unclonable Functions (PUFs). However, the state of the art still often relies on microelectronic integration to achieve security by obscurity insted of more fundamental security guarantees. While effective, system-level tamper protection is only used in few devices such as Hardware Security Modules (HSMs) and card payment terminals. Due to the high cost and low performance of HSMs in particular, they remain relegated to niche applications such as Transport Layer Security (TLS) certificate issuance and payment data processing.

In this thesis, Jan Sebastian Götte introduces the Inertial Hardware Security Module (IHSM), a new architecture for low-cost hardware security modules that provide high-level active tamper protection, while supporting computing payloads of much larger size, weight and power dissipation compared to conventional HSMs. In an IHSM, the costly and difficult to source tamper-sensing mesh of a conventional HSM is replaced by a mesh made from simple PCBs that is rotating at high speed around the payload. Since the mesh is rotating, it cannot be manipulated, and the security of conventional meshes created in bespoke manufacturing processes can be achieved using much simpler and less expensive construction techniques. We present the results of a survey of approximately 30 real world tamper sensing mesh implementations. We deduce design criteria for secure meshes and contextualize our design. We further motivate the necessity of secure hardware by presenting an analysis of problematic aspects in the hardware security design of Germany's new national electronic health record system.

To pave the way for practical implementations of IHSM technology, we present solutions to key engineering challenges in IHSM construction.

We present a design and analysis of highly symmetric planar inductors for rotating wireless power transfer. We present a high-fidelity, low-cost monitoring system for security meshes that is based on the principles of Time-Domain Reflectometry (TDR). We validate our system and find that it is able to reliably detect several classes of advanced physical attacks. We find that our system is sensitive enough to detect differences between identical copies of the same mesh, suggesting PUF-like properties.

Applying IHSM technology, we analyse two use cases that are unlocked by the increased size and power dissipation capability of IHSMs. In the first analysis, an IHSM-secured relay node for Quantum Key Distribution (QKD) systems is proposed, enabling their practical implementation across arbitrary distances, which requires trusted relay stations due to fundamental physical limitations. In the study, IHSMs are adapted for such high-security QKD relays by securing the IHSM mesh passthrough with a secondary tamper-sensing mesh. In this setup, a bracket design is proposed that supports passing through optical fibers at low loss.

The second proposed use case adapts an IHSM enclosure to the size, power and thermal dissipation requirements of a high-power server to support co-located secure Multiparty Computation (MPC) workloads. In practical MPC deployments, nodes are distributed across data centers to avoid a single point of failure for physical attacks. As a result, practical MPC deployments are limited by network bandwidth and latency constraints. Using IHSMs, physically secured MPC nodes can be deployed within the same data center, increasing bandwidth, reducing latency and unlocking a new performance spectrum.

Use of Artificial Intelligence in This Thesis

This thesis has been written during the years of 2020 - 2025. In this time, Artificial Intelligence (AI) technology including Large Language Models (LLMs) has entered widespread adoption. I have used such LLM systems in the preparation of this thesis. At the time this thesis was written, LLMs were a powerful and useful technology, but often produced wrong output. Thus I used the following list of observations to guide my LLM use during the writing of this thesis.

1. Passing text through an LLM is an imprecise operation. Especially when large amounts of text are passed through an LLM, despite clear instructions such as "only fix spelling errors," the LLM output might deviate from the source text. Therefore, the document text should never be passed through the LLM, and the LLM should be prompted to point out problems, or to produce a list of suggestions for improvements instead.

2. Contemporary LLMs are bad at summarizing text that contains novel concepts. LLM summaries of text often converge to a re-stating of the general consensus on the text's main topic. Where the source text deviates from conventional wisdom or makes novel points, an LLM summary will likely mis-represent those conclusions. Additionally, LLMs are bad at capturing the point of a text. Unless extreme care is taken when prompting, it is easy to lead an LLM to produce an inaccurate summary of a text that agrees with the prompt, but misses the gist of the text. Therefore, extreme caution should be applied when using an LLM for summarization, and LLM output should be checked diligently in such instances.

3. Contemporary LLMs are bad at generating text from scratch. Espe-

cially on topics of academic interest that are novel and that do not have well-known answers that can be found in the training corpus for these models, in general they will not produce useful text when prompted. Therefore, LLMs should never be used to generate novel text.

4. Contemporary LLMs are bad at giving references. Prompts that ask for academic references on a topic are likely to produce non-existing "hallucinated" references. The existing references an LLM is most likely to dig up usually occur on the first page of a web search on the topic, or are frequently cited in literature on the topic. Thus, LLMs should never be directly queried for references. When researching a new concept, a better use of an LLM is the generation of query strings for search engines like Google Scholar.

Applying these observations, I never copied text from the LLM into this thesis. Where I edited the text of this thesis using suggestions from LLM output, I critically evaluated the LLM output and carefully considered each edit. Following are some examples of how I used LLMs in the writing of this thesis.

For checking spelling and grammar, the LLM was prompted with an instruction to review the text and output a list of errors. The list was then reviewed and the errors were fixed in the source document by hand. An example prompt for the LLM in this case might be: "The attached file contains the LaTeX source code of a chapter of an doctoral thesis titled '...'. Review the text and list any mistakes in spelling or grammar."

For improving formulation patterns, the LLM was prompted with a short excerpt of text of at most two paragraphs and instructions asking for an improved version of the text. In response to such a prompt, the LLM will often change the meaning of parts of the text. Thus, I used the output as a reference example, and manually adjusted the source document applying parts of the LLM response where fitting. An example prompt in this case might be: "The following text are two paragraphs from a chapter on '...' in a PhD thesis on '...' . Improve the wording of these paragraphs to make them easier to read and understand.".

For improving the structure of the text, the LLM was prompted with an instruction to review the text and output a list of recommendations. The list was then reviewed, and changes were made to the source document by hand. An example prompt in this case might be: "The attached document contains the LaTeX source code of a chapter of a PhD thesis on '...' . Critically assess the structure and organization of the chapter and write a list of suggestions for improvement."

In accordance with the recommendations of the University and State Library Darmstadt regarding the labelling and documentation of AI-generated materials dated September 22, 2025[W217], instances where I used an LLM to edit parts of the text of this thesis as described above have not been explicitly labelled in the text. The LLM in this use assumes a similar role a human editor might assume reviewing the text.

Besides the use of LLMs as described above, a specialized machine translation tool was used to create the German translation of the abstract at the beginning of this thesis. This use is marked explicitly.

Web sources

[W217] *Recommendations of the University and State Library Darmstadt for Labelling and Documenting AI-generated Content*. 2025-09-22. URL: `https://www.ulb.tu-darmstadt.de/ki-doku` (visited on 2025-10-24) (cit. on p. ix).

Contents

Chapter 1

Introduction

> It's not for lack of ideas or possibilities. It's that we actually have to start taking seriously the shifts that are going to be required to do this thing—to build tech that rejects surveillance and centralized control—whose necessity is now obvious to everyone.
>
> – *Meredith Whittaker* [W97]

Contents

No Gods, No Masters is an anarchist slogan originating in the 19th century that expresses a rejection of authorities [35, 99, 27]. In modern cryptography, it is generally seen as best practice to have the least amount of parties possible involved in any computation. Most cryptographic problems are easily solved by involving a trusted third party (TTP). Yet, cryptographers have time and again vocally rejected attempts to involve third parties in cryptographic protocols [4, 5, 14, 220].

Considerable research has been focused on creating a versatile set of tools to perform tasks as diverse as secure communication [9, 165, 57, 227], oblivious database access [46, 7, 218], and even general computation [93, 17, 46] in a decentralized way that avoids trusted authorities. While politically, this blanket rejection of authority represents a fringe viewpoint, in cryptography it has a long tradition originating with the Cypherpunk and Hacker movements [12, W114, 127, W164], and extending throughout mainstream academic cryptography.

While the aforementioned cryptographic tools enable a large gamut of use cases in theory, in practice cryptographic systems are still routinely compromised [81, 91, 231, 212, 163, W209, W200]. A fundamental flaw of any practical cryptographic system is that secure algorithms have to run on hardware, and even today, average computing hardware provides little physical security [96, 153, 140, 174]. *Hardware Security Modules* are a class of devices specifically designed to execute cryptographic algorithms while providing strict physical security guarantees, but these systems are expensive, and their physical security is often questionable [195, 14], which we wi elaborate further in Chapter 3. Anderson [14] writes on HSMs and their security standards:

> Security economics remains a big soft spot, with security chips being in many ways a market for lemons. A banker buying HSMs probably won't be aware of the huge gap between FIPS[1] level 3 and level 4, and understand that level 3 can sometimes be defeated with a Swiss army knife. The buying incentive there is compliance, and where real security clashes with operations it's not surprising to see weaker standards designed to make compliance easier.
>
> *Anderson [14] p. 629*

[1] Anderson here refers to the US national HSM security standard FIPS 140 [1, 2]

In this thesis, we aim to fill this gap in easily obtainable, secure hardware and extend the level of protection afforded by cryptographic protocol design down the technology stack to the hardware level. We propose a new HSM design that unlike existing designs can be manufactured at low cost and without access to specialized tools.

We publish our design fully open source, including all detials necessary for replication. A fundamental principle in cryptographic engineering is Kerckhoffs' principle[2], named after Dutch military cryptographer Auguste Kerckhoffs. Kerckhoffs' principle expresses that the security of a cryptographic system should only depend on the secrecy of its keys, not on the secrecy of its design. Existing commercial designs routinely contravene Kerckhoff's principle by applying the widespread industry practice of *Security by Obscurity*. Even in academic related work, the principle is sometimes violated by omitting implementation and methodological details in the interest of patents and commercial exploitation. By publishing all details of our research into HSMs and their components, we provide the foundation for future independent research.

Beyond applying Kerckhoffs' principle, publishing our design also enables independent replication. Our design is based entirely on standard components and does not require bespoke manufacturing processes. Both commercial and academic existing HSM tamper sensing designs require bespoke manufacturing methods or custom integrated circuits (ICs) [P196, 118, 80, 116]. Custom ICs require a large up-front financial commitment to produce. Bespoke manufacturing methods may require custom machines, training, and specialty materials, also incurring a high startup cost. This creates a single point of failure in the manufacturer, and opens up an opportunity for a hardware supply-chain attack [102]. Such supply chain attacks can be mitigated by independently manufacturing our design.

1 A Note on Hardware Security Module Terminology

In this thesis, we use the term *Hardware Security Module (HSM)* to refer to a security device that has the following three properties.

1. A HSM targets the prevention of any conceivable physical attack. In

[2] Petitcolas [W206] contains a high-quality OCR'ed copy of the original source, as well as a translation of the cited part from French. The original source is Kerckhoffs [133].

particular, this includes intrusion attempts such as careful drilling or cutting into the device from any direction.

2. A HSM includes tamper sensors that when triggered result in an active tamper response, usually deleting all cryptographic secrets and rendering the device inoperable.

3. A HSM's tamper sensing and response subsystem is continuously powered from a backup power supply, usually a battery. Loss of power triggers the tamper response.

This use of the term *HSM* aligns with common usage of the term both in the academic literature and in everyday conversation. Particularly the requirement of active tamper detection and response is crucial to distinguish a HSM from simpler devices such as TPMs, smart cards or secure enclaves in SoCs. Note that our use of the term HSM is slightly different from its use in government standards, from its use in the PCI SSC (Payment Card Industry Security Standards Council) standards, and from its industry use.

In industry, the term HSM is often used for solutions that are only logically segregated and that do not include any particular defense against hardware attacks. Our conjecture is that this is a consequence of the standardization landscape, where for applications outside of card payment processing the US FIPS 140-22 [1] standard was central to the industry. Despite encompassing both devices that include active tamper detection and response, FIPS 140-2 did not draw a distinction in its terminology between the two classes.

1.1 Use in government standards

Under the still widely used US national standard FIPS 140 in in its 2002 version 2 [1], a HSM would be called a *Multiple-Chip Cryptographic Module* that conforms to the standard's *Security Level* 4 out of 4. Interesting to note are that only level 4 requires any active tamper detection and response, so devices compliant only up to levels 3 and below do not align with our HSM definition. Futher of note is that according to the standard, a single-chip solution does not require any tamper detection and response either to meet the standard's security level 4, which is in misalignment with our definition. The standard's 2019 updated version FIPS 140-3 [2] defers to the international standards ISO/IEC 19790 and 24759.

ISO/IEC 19790 [W125] and ISO/IEC 24759 [W126] call what we call a HSM a *Hardware Cryptographic Module* corresponding with the standards *Security Level 4*. However, these standards only require active tamper detection and response when cryptographic secrets are transmitted in plaintext between chips.

1.2 Use in card payment processing (PCI SSC) standards

The Payment Card Industry Security Standards Council (PCI SSC) is an association of credit card network operators that defines standards for all layers of card payment processing, from card payment terminals in stores to the handling of payment data in online shop backend systems.

PCI SSC terminology aligns with our definition and with common everyday use of the term HSM. In PCI SSC terminology, a HSM is a crytographic device that has active tamper detecion and response circuitry. However, PCI SSC terminology differs from our use of the term HSM in one nuance: In PCI SSC terminology, a HSM is specifically a datacenter device used for backend processing of payment data. The general class of "hardware devices performing some security function with or without particular physical security requirements" that ISO/IEC 19790 and other standards call a *Hardware Cryptographic Module*, in PCI SSC terminology is termed *Secure Cryptographic Device (SCD)* in more recent standard versions, which was updated from the previous term *Tamper-Resistant Security Module (TRSM)*. Other than HSMs, PCI SSC includes smartcards and card payment terminals in this category. Card payment terminals, referred to as *Pin-Entry Device (PED)* in PCI SSC standards, have to include a surprising amount of active tamper detection and response functionality including partial coverage of areas like their main cryptographic processor and smart card reader by battery-backed tamper-sensing meshes. Under our definition, these devices can be classified as a type of HSM.

1.3 Tamper-Sensing Meshes

In this thesis, we use the terms *Tamper-Sensing Mesh* and *Security Mesh* synonymous. We use both terms to refer to any electrical circuit whose path is laid out to cover a surface with the intent of detecting attempts at drilling, cutting or otherwise manipulating this surface. While the term *Security Mesh* is more concise, it is less clear to people unfamiliar with the matter. It is also polysemous, and depending on context can also refer to

woven or stamped metal meshes used as fences or as screens in front of windows to prevent break-ins. As a result, it is harder to use in online searches, and when using Large Language Models (LLMs), it frequently leads to amusing hallucinations.

2 Inertial Hardware Security Modules

In this thesis, we propose Inertial Hardware Security Modules (IHSMs) to fill the gap of protecting systems that handle highly sensitive data but that cannot use conventional HSMs for cost or performance reasons. In a system with a secure software stack, the role of a HSM is to secure the hardware part of the stack. The basic approach of a HSM is to combine a secure software stack with tamper sensors connected to a fast self-destruct mechanism. The tamper sensors are tasked with detecting any physical attack an attacker could mount on the device. Common classes of such sensors include environmental sensors such as temperature or radiation sensors that detect attempts at causing controllable faults in the HSM by heating, cooling or irradiating it. Building on the basic protection offered by such sensors, *tamper-sensing meshes* are often employed. These *meshes* are flexible foils containing circuit traces that are attached to the HSM's enclosure to detect attempts at penetrating the shell of the device with probes. Tamper-sensing meshes usually are the primary line of defense against most physical attacks. They are very effective at mitigating a large variety of physical attacks, but they are difficult to construct securely as they usually require bespoke manufacturing processes. As a result, they are currently only used in niche applications, and even there not every realization is equally secure. The self-destruct mechanism can be hardware or software that quickly and securely destroys all cryptographic secrets, thereby rendering the device worthless to an attacker.

IHSMs are a new design approach that utilizes mechanical motion to create secure tamper-sensing meshes from simple components. IHSMs solve the issue of creating an impenetrable tamper-sensing envelope by replacing the bespoke tamper-sensing mesh foil with a set of simple, rigid meshes made from commodity Printed Circuit Boards (PCBs) that are rotating at high speed. In motion, these simple PCB tamper-sensing meshes are as secure as the much more sophisticated bespoke foils used in conventional HSMs against an attacker with access to commercially available tools, yet

they are simpler and less expensive to manufacture. To verify that the mesh is rotating correctly, an accelerometer is placed on the rotating mesh, and its centrifugal force reading is used to validate its path of motion.

IHSMs enable the protection of much larger payloads compared to conventional mesh designs, and they can support larger power dissipation. Combined with their low cost, this enables the implementation of high-level hardware security in applications that previously would not have been possible to secure.

To the best of our knowledge, IHSMs are the first fully open source, replicable HSM with advanced tamper sensing features. Across application domains, IHSMs can be applied to gain resistance to physical attacks in scenarios where conventional HSMs were not used because of cost, computing power or implementation effort. Where conventional HSMs come as fully integrated devices that only expose limited APIs to their users, IHSMs at their core are just an enclosure that the user can put whatever hardware they need into, adapting the tamper response to their application's needs. Since the simpler tamper-sensing mesh construction of IHSMs scales to larger payload volumes, entire servers can be protected—something that is impossible with conventional HSMs. Since the mesh in an IHSM is constantly moving, unlike a mesh in a conventional HSM, it does not have to entirely cover the payload. Instead, it can have gaps that allow for air flow between outside and inside, enabling active cooling of the IHSM's payload. This cooling capability increases computing power by increasing feasible payload power dissipation by orders of magnitude [142].

3 Research Questions and Contributions

Based on the current state of the field of hardware security, we deduce three overarching research questions for this thesis that progress from theory to practical deployment.

1. What is the state of the art in commercial tamper sensing mesh implementations?

2. What are criteria and approaches for the design of secure tamper sensing meshes?

3. Can we achieve physical security without relying on a conventional

tamper-sensing meshes that requires a bespoke manufacturing process?

4. Can we monitor tamper-sensing meshes at a higher detail level than the state of the art of a single, scalar measurement?

5. Can we improve the ripple voltage performance of Wireless Power Transfer (WPT) through rotating joints to adapt it to IHSM applications?

6. What applications does our IHSM technology open up through its increase in power dissipation and size capabilities?

We answer our first research question in two parts. In Chapter 2, we analyze the hardware security design of Germany's new national electronic health record system. Our analysis unveils a combination of problematic choices resulting from conflicting constraints and lack of awareness. In Chapter 3, we present the results of a survey across approximately 30 real world tamper sensing mesh implementations, analyzing common design features.

The latter half of our survey in Chapter 3 answers our second research quesion. From our analysis of this large corpus of devices, we deduce a list of design criteria that can be applied to increase the security of any tamper sensing mesh implementation.

To answer our third research question, in Chapter 4 we propose the Inertial Hardware Security Module (IHSM), a new type of HSM that extends the high level of protection offered by the modern cryptographic software stack down to the hardware level, enabling secure computation in insecure places. IHSMs can be built from basic, off-the-shelf components and do not require bespoke manufacturing processes.

IHSMs come with unique power supply constraints since their rotating mesh must be continuously powered. A straightforward solution utilizes Wireless Power Transfer using planar inductors, but existing WPT designs exhbit a ripple voltage due to an asymmetry of conventional planar inductors. This leads to our fourth research question, which we solve in Chapter 6 with the design and experimental evaluation of a new, generalized class of *twisted* planar inductors that reduces voltage ripple in rotating shaft setups.

To answer our fifth research question, in Chapter 5 we propose improvements to the state of the art in HSM tamper sensors based on the use of

low-cost, embeddable Time-Domain Reflectometry (TDR). Our improvements can be applied to both IHSMs and conventional HSMs.

Finally, we answer our last research question by showing in two case studies how an end-to-end design of an IHSM-secured data processing system could look like. Both case studies concern scenarios that IHSMs unlock that were previously infeasible using conventional HSMs: In Chapter 7, we explore how IHSMs enable long-range Quantum Key Distribution (QKD) networks using trustable physically secured relay nodes and in Chapter 8 we elaborate how datacenter-scale Secure Multiparty Computation (SMPC) clusters can be created using IHSM enclosures with commercial server hardware.

4 Contributions

Through this thesis, we make contributions advancing the state of hardware securty across several related sub-fields. Our contributions include:

1. We conduct the first large-scale survey of tamper sensing measures in the real world, analyzing approximately 30 devices.

2. From our real world observations, we systematize tamper sensing mesh construction techniques and we provide a list of criteria improving mesh security.

3. We experimentally analyze the impact of Computed Tomography (CT) imaging on mesh security.

4. We propose the IHSM, a new concept for HSM design based on a rotating mesh that increases payload size and power dissipation capacity while simultaneously allowing for simpler meshes constructed from standard components.

5. We show experimental results on IHSM mesh performance obtained with a prototype IHSM.

6. We introduce an algorithm for the automatic layout of tamper-sensing meshes and its implementation on top of a popular, open-source Electronic Design Automation (EDA) tool.

7. We introduce a high-fidelity mesh monitoring approach that uses Time-Domain Reflectometry (TDR).

8. We show a low-cost implementation of our TDR monitoring approach.

9. We evaluate the performance of our TDR monitoring implementation and demonstrate its response to a large set of attacks. We show that it reliably distinguishes identical copies of the same mesh specimen, suggesting PUF-like behavior.

10. We introduce a generalized design approach for low-loss planar inductors that out-peform prior approaches in parasitic capacitance, self-resonant frequency and rotational symmetry.

11. We apply our design approach to the problem of Wireless Power Transfer to the rotating mesh of an IHSM.

12. We conduct an exhaustive experimental evaluation of the rotational symmetry of a large set of planar WPT inductors created using our approach.

13. We analyze physically secure Quantum Key Distribution relays as an IHSM use case and develop a low-loss fiber optic passthrough that supports an additional, secondary, independently rotating mesh shielding the shaft passthrough of the IHSM's primary mesh.

14. We explore IHSMs for co-located high performance Multiparty Computation (MPC) setups. We demonstrate a fan-driven IHSM mesh concept for high-availability scenarios that removes motors as a single point of failure while providing sufficient airflow for cooling high-power server components.

We chose to publish all of our research as open source and unencumbered by patents to enable widespread adoption. IHSMs can be custom built with only basic manufacturing capabilities at small scale and enable the deployment of secure computation in insecure places even to small organizations such as university research departments, NGOs and small businesses.

Looking at the practice of applied hardware security, we observe that despite ample availability of commercial solutions promising easy hardware security, clearly there is still a lack of solutions that provide the adaptability necessary for some real use cases at low enough cost. By publishing the tamper-sensing technology we developed during the making of this thesis as open source hardware designs, we aim to provide this missing

building block to provide high-level hardware security in real-world applications. Our hardware designs can be adapted to devices ranging from Single-Board Computers (SBCs) to servers, they are compatible with non-computing applications like Quantum Key Distribution (QKD) and their design approaches can even be integrated into existing HSM designs to provide better security at little additional cost.

Web sources

[W97] Andy Greenberg. *Signal Is More Than Encrypted Messaging. Under Meredith Whittaker, It's Out to Prove Surveillance Capitalism Wrong.* WIRED Magazine. 2024-08-28. URL: https://www.wired.com/story/meredith-whittaker-signal/ (visited on 2025-06-13) (cit. on p. 1).

[W114] Eric Hughes. *A Cypherpunk's Manifesto.* URL: https://www.activism.net/cypherpunk/manifesto.html (visited on 2025-11-18) (cit. on p. 2).

[W125] *ISO/IEC 19790:2025.* ISO. URL: https://www.iso.org/standard/82423.html (visited on 2025-05-15) (cit. on pp. 5, 24, 38).

[W126] *ISO/IEC 24759:2025.* ISO. URL: https://www.iso.org/standard/82424.html (visited on 2025-04-08) (cit. on pp. 5, 24, 115, 119).

[W164] Moxie Marlinspike. *We Should All Have Something To Hide.* Blog of Moxie Marlinspike. 2013-06-12. URL: https://moxie.org/2013/06/12/we-should-all-have-something-to-hide.html (visited on 2025-11-18) (cit. on pp. 2, 217).

[W200] *Pakistan: Mass Surveillance and Censorship Machine Is Fueled by Chinese, European, Emirati and North American Companies.* Amnesty International Security Lab. 2025-09-09. URL: https://securitylab.amnesty.org/latest/2025/09/pakistan-mass-surveillance-and-censorship-machine-is-fueled-by-chinese-european-emirati-and-north-american-companies/ (visited on 2025-11-27) (cit. on p. 2).

[W206] Fabien Petitcolas. *Kerckhoffs' Principles from « La Cryptographie Militaire ».* The information hiding homepage. URL: http://www.petitcolas.net/steganography/ (visited on 2025-11-18) (cit. on p. 3).

[W209] *Predator Files: Technical Deep-Dive into Intellexa Alliance's Surveillance Products*. Amnesty International Security Lab. 2023-10-06. URL: https://securitylab.amnesty.org/latest/2023/10/technical-deep-dive-into-intellexa-alliance-surveillance-products/ (visited on 2025-11-27) (cit. on p. 2).

Patent References

[P196] Johannes Obermaier, Vincent Immler, and Robert Hesselbarth. "PUF-film and Method for Producing the Same". U.S. pat. 11586780B2. Fraunhofer Gesellschaft zur Foerderung der Angewandten Forschung eV. 2023-02-21 (cit. on pp. 3, 37).

References

[1] (US) National Institute of Standards and Technology. *Security Requirements for Cryptographic Modules*. Federal Information Processing Standard (FIPS) 140-2. U.S. Department of Commerce, 2002-12-03. DOI: 10.6028/NIST.FIPS.140-2 (cit. on pp. 2, 4, 24, 115, 119).

[2] (US) National Institute of Standards and Technology. *Security Requirements for Cryptographic Modules*. Federal Information Processing Standard (FIPS) 140-3. U.S. Department of Commerce, 2019-03-22. DOI: 10.6028/NIST.FIPS.140-3 (cit. on pp. 2, 4, 24, 38, 66).

[4] Hal Abelson et al. "The Risks of Key Recovery, Key Escrow, and Trusted Third-Party Encryption". In: *World Wide Web J.* 2.3 (1997-06-01), pp. 241–257. ISSN: 1085-2301 (cit. on pp. 2, 22).

[5] Harold Abelson et al. "Keys under Doormats". In: *Commun. ACM* 58.10 (2015-09-28), pp. 24–26. DOI: 10.1145/2814825 (cit. on pp. 2, 22).

[7] Carlos Aguilar-Melchor et al. "XPIR : Private Information Retrieval for Everyone". In: *Proceedings on Privacy Enhancing Technologies* 2016.2 (2016-04-01), pp. 155–174. DOI: 10.1515/popets-2016-0010 (cit. on p. 2).

[9] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. "The Double Ratchet: Security Notions, Proofs, and Modularization for the Signal Protocol". In: *Advances in Cryptology – EUROCRYPT 2019*. Ed. by Yuval Ishai and Vincent Rijmen. Vol. 11476. Cham: Springer International Publishing, 2019, pp. 129–158. DOI: 10.1007/978-3-030-17653-2_5 (cit. on p. 2).

[12] Patrick D. Anderson. *Cypherpunk Ethics: Radical Ethics for the Digital Age*. London: Routledge, 2022-04-24. 142 pp. DOI: 10.4324/9781003220534 (cit. on p. 2).

[14] Ross Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 1st ed. Wiley, 2020-12-22. DOI: 10.1002/9781119644682 (cit. on pp. 2, 22, 24, 36–38, 43, 64, 75, 77, 78, 92, 118).

[17] Yonatan Aumann and Yehuda Lindell. "Security Against Covert Adversaries: Efficient Protocols for Realistic Adversaries". In: *Journal of Cryptology* 23.2 (2010-04), pp. 281–343. DOI: 10.1007/s00145-009-9040-7 (cit. on pp. 2, 221).

[27] Tashina Blom. "No Gods No Masters: Anarchist Mots de Mémoire from Titles to T-Shirts". In: Sophie Van Den Elzen and Ann Rigney. *Memory and the Language of Contention*. BRILL, 2025-03-10, pp. 231–247. DOI: 10.1163/9789004692978 (cit. on p. 2).

[35] Romain Broussais. "Les Origines de La Devise Anarchiste « Ni Dieu Ni Maître » : Une Généalogie Discutable". In: *Histoire Politique* 46 (2022-02-01). DOI: 10.4000/histoirepolitique.2452 (cit. on p. 2).

[46] Benny Chor, Oded Goldreich, and Eyal Kushilevitz. "Private Information Retrieval". In: () (cit. on p. 2).

[57] Benjamin Dowling, Paul Rösler, and Jörg Schwenk. "Flexible Authenticated and Confidential Channel Establishment (fACCE): Analyzing the Noise Protocol Framework". In: *Public-Key Cryptography – PKC 2020*. Ed. by Aggelos Kiayias et al. Vol. 12110. Cham: Springer International Publishing, 2020, pp. 341–373. DOI: 10.1007/978-3-030-45374-9_12 (cit. on p. 2).

[80] Kathrin A Garb. "Tamper-Sensitive Design of PUF-Based Security Enclosures" (cit. on pp. 3, 115, 118, 124, 198).

[81] Barton Gellman and Ashkan Soltani. "NSA Infiltrates Links to Yahoo, Google Data Centers Worldwide, Snowden Documents Say". In: *The Washington Post* (2013-10-30). ISSN: 0190-8286 (cit. on p. 2).

[91] Adam Goldman. "'Unrestrained' Chinese Cyberattackers May Have Stolen Data From Almost Every American". In: *The New York Times. World* (2025-09-04). ISSN: 0362-4331 (cit. on p. 2).

[93] Gerhard Goos et al. "Information Theoretically Secure Communication in the Limited Storage Space Model". In: *Advances in Cryptology — CRYPTO' 99*. Ed. by Michael Wiener. Vol. 1666. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 65–79. DOI: 10.1007/3-540-48405-1_5 (cit. on p. 2).

[96] Johannes Götzfried et al. "Cache Attacks on Intel SGX". In: *Proceedings of the 10th European Workshop on Systems Security*. EuroSec'17. New York, NY, USA: Association for Computing Machinery, 2017-04-23, pp. 1–6. DOI: 10.1145/3065913.3065915 (cit. on p. 2).

[99] Daniel Guérin. *No Gods No Masters: An Anthology of Anarchism*. Trans. by Paul Sharkey. Complete unabridged ed. Edinburgh, Scotland Oakland, CA: AK Press, 2005. 699 pp. ISBN: 978-1-904859-25-3 (cit. on p. 2).

[102] Jacob Harrison, Nathan Jessurun, and Mark Tehranipoor. "SoK: A Security Architect's View of Printed Circuit Board Attacks". In: 34th USENIX Security Symposium (USENIX Security 25). 2025, pp. 1907–1924. ISBN: 978-1-939133-52-6 (cit. on p. 3).

[116] 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)*. 2018 IEEE International Symposium on Hardware Oriented Security and Trust (HOST). 2018-04, pp. 49–56. DOI: 10.1109/HST.2018.8383890 (cit. on pp. 3, 37, 115, 116, 118, 124).

[118] Vincent Immler et al. "Secure Physical Enclosures from Covers with Tamper-Resistance". In: *IACR transactions on cryptographic hardware and embedded systems.* (2019). DOI: 10.13154/tches.v2019.i1.51-96 (cit. on pp. 3, 77–79, 82).

[127] Craig Jarvis. *Crypto Wars: The Fight for Privacy in the Digital Age: A Political History of Digital Encryption*. 1st ed. CRC Press, 2020-12-14. ISBN: 978-1-00-312367-5 (cit. on pp. 2, 22).

[133] Auguste Kerckhoffs. "La Cryptographie Militaire". In: Journal Des Sciences Militaires 9 (1883-01), pp. 5–38 (cit. on p. 3).

[140] Paul Kocher et al. "Spectre Attacks: Exploiting Speculative Execution". In: *Communications of the ACM* 63.7 (2020), pp. 93–101. DOI: 10.1145/3399742 (cit. on p. 2).

[142] Tony Kordyban. *Hot Air Rises and Heat Sinks: Everything You Know about Cooling Electronics Is Wrong*. ASME, 1998. ISBN: 978-0-7918-0074-4 (cit. on pp. 7, 86).

[153] Moritz Lipp et al. "Meltdown: Reading Kernel Memory from User Space". In: *Communications of the ACM* 63.6 (2018), pp. 46–56. DOI: http://dx.doi.org/10.1145/3357033 (cit. on p. 2).

[163] Bill Marczak and John Scott-Railton. *Graphite Caught: First Forensic Confirmation of Paragon's iOS Mercenary Spyware Finds Journalists Targeted*. Citizen Lab, University of Toronto, 2025-06-12 (cit. on p. 2).

[165] Moxie Marlinspike and Rolfe Schmidt. *The Double Ratchet Algorithm*. Ed. by Trevor Perrin. 2025-11-04 (cit. on p. 2).

[174] Daniel Moghimi et al. "TPM-FAIL: TPM Meets Timing and Lattice Attacks". In: *Proceedings of the 29th USENIX Security Symposium*. USENIX Security Symposium. USENIX Association, 2020-08, pp. 2057–2073. ISBN: 978-1-939133-17-5 (cit. on p. 2).

[195] Johannes Obermaier and Vincent Immler. "The Past, Present, and Future of Physical Security Enclosures: From Battery-Backed Monitoring to PUF-based Inherent Security and Beyond". In: *Journal of Hardware and Systems Security* 2 (2018), pp. 289–296. DOI: 10.1007/s41635-018-0045-2 (cit. on pp. 2, 37, 63, 78, 218).

[212] Cooper Quintin, Rebekah Brown, and John Scott-Railton. *Something to Remember Us By: Device Confiscated by Russian Authorities Returned with Monokle-Type Spyware Installed*. Citizen Lab, University of Toronto, 2024-12-05 (cit. on p. 2).

[218] Leonie Reichert et al. “Menhir: An Oblivious Database with Protection against Access and Volume Pattern Leakage”. In: *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security*. ASIA CCS ’24. New York, NY, USA: Association for Computing Machinery, 2024-07-01, pp. 1675–1690. DOI: 10.1145/3634737.3657005 (cit. on p. 2).

[220] Phillip Rogaway. “The Moral Character of Cryptographic Work”. In: *Advances in Cryptology*. ASIACRYPT 2015. Vol. 9452 & 9453. LNCS. Auckland, New Zealand: Springer, 2015, p. XVIII. DOI: 10.1007/978-3-662-48800-3 (cit. on pp. 2, 22, 217).

[227] Sajin Sasy and Ian Goldberg. “SoK: Metadata-Protecting Communication Systems”. In: *Proceedings on Privacy Enhancing Technologies* 2024.1 (2024-01), pp. 509–524. DOI: 10.56553/popets-2024-0030 (cit. on p. 2).

[231] John Scott-Railton et al. *By Whose Authority? Pegasus Targeting of Russian & Belarusian-speaking Opposition Activists and Independent Media in Europe*. Citizen Lab, University of Toronto, 2024-05-30 (cit. on p. 2).

Chapter 2

The German ePA: A Motivating Counter-Example

> The most dangerous phrase in the language is "We've always done it this way!".
>
> – *attributed to Grace Hopper[W277, W213]*

Contents

To do
FIXME: Proper citation here

This part is based on a short paper written by me and presented by Jan Sebastian Götte at the HS3 workshop at ESORICS 2025.

Looking at the landscape of computer security solutions, we are presented with a wide variety of vendors and products that may give the impression that hardware security is a solved problem. Vendors sell various claims rangning from *"You don't need hardware security, just do it in the cloud!"* [W261, W172, W115, W10, W92, W273] to *"Buy our HSM and you will be secure!"* [W260, W248]. In practice, things are not as easy and even well-intentioned projects still often go awry on the hardware security dimension. To motivate our research into physical security in this thesis, in this chapter we will have a look at one such project that was done by capable people with the best intentions, yet it resulted in a hardware security design that is dangerously inadequate for the purpose.

Beginning May 2025, after several delays, Germany has started the nation-scale rollout of its new electronic medical record system, named ePA (short for *elektronische Patientenakte*, "electronic patient record") [W139]. The system aims to create a national database accessible to all healthcare providers that holds the complete electronic medical records of all publically insured people living in Germany. The system aims to replace paper-based workflows that are error-prone and lead to healthcare providers often only having access to a subset of patient's medical records. Data in scope for the system includes medical letters, laboratory results, and medical imaging files.

Due to Germany's mandatory health insurance laws, the system's user base encompasses the majority of all German residents, approximately 90%. People who have replaced their public health insurance with private insurance as of now are not subject to the system. In Germany, by law private health insurance is only available to people from the top 10th percentile of household income. This means that the system disproportionally affects people who have low income, creating an equity issue. While it is possible to opt out from the use of the new digital record, the process of opting out is difficult. Additionally, the government and health insurance providers have publically depicted the system in a one-sidedly positive way, meaning that it is unlikely the majority of people subject to the system have a comprehensive understanding of the system's benefits and risks that would be necessary for an informed decision.

While there has been loud criticism of the system's security from civil society organizations such as digital rights nonprofit organization Chaos Computer Club (CCC) [W138] and several severe security flaws have been

demonstrated practically, this criticism has largely been ignored by the political structures in charge. We observe that despite this civil society outrage and the system's large scale, it has received little attention from the academic cryptography and information security community.

In this chapter, we aim to highlight some unconventional cryptographic engineering decisions in the system. In particular, we point out that the system's core per-user secrets are kept in a rudimentary key escrow system whose security is based on engineering assumptions, not on cryptographic principles. Furthermore, we observe that by specification, the individual user keys of the system are derived from a per-user cleartext salt based on a system-wide long-term secret with only 256 bits of entropy[1]. Finally, we note that according to specification, the only physical security requirement for the protection of this highly sensitive secret is a "hard, opaque potting material", with no tamper detection and response required.

We base our analysis of the ePA on the system's publicly available standards in their latest version as of the writing of the paper underlying this chapter in April 2025, describing version 3.0 of the healthcare record system [[W]83, [W]85]. We note that hypothetically, the implementation might deviate from these standards and be more secure. The reference implementation provided by the specification authority [[W]90] follows the specified minimum requirements closely. As of now, there is no meaningful way for either the public or for researchers such as us to ascertain the concrete implementation security of the system.

1 The Design of ePA

ePA is embedded into Germany's national public healthcare backend system "Telematikinfrastruktur" (abbreviated TI; German for "telematics infrastructure"). TI is a highly complex system, and a detailed description would exceed the limits of this analysis. Briefly put, TI consists of a shared demilitarized zone (DMZ) that parties like insurance providers and healthcare providers connect to through a VPN. At the client location, usually an individual doctor's office or a hospital, this VPN connection is terminated by a specialized VPN appliance named "Konnektor" that simultaneously

[1] In previous versions of the standard [[W]84, [W]86], there were two escrow services, with both keys used in layers to reduce the risk of a compromise of either one. The current standard only requires one escrow service, and drops the entropy requirement of the root keys from 512 bits to 256 bits. The apparent reason for the long-term nature of these keys is that they are updated manually.

acts as a trusted component inside the client network hosting some software for purposes such as authentication. The Konnektor contains several smart cards that store keys used for authentication. Konnektor devices are offered by several vendors and healthcare providers like doctor's offices are indivudally responsible for purchasing and maintaining a Konnektor.

Every person enrolled in the system as well as every healthcare professional providing services under it is issued an ID card that contains a smart card with keys to authenticate towards the central infrastructure. The primary use of these smart cards previously was to automatically provide personal information such as name, birth date, address and insurance enrollment status when an enrolled person visits a healthcare provider.

ePA is implemented inside the TI system. Its centralized services are accessed by healthcare providers through the TI's VPN, and by patients through proxy servers connected to TI's VPN. Patient records are encrypted and decrypted inside TI's backend systems. Smart cards authenticate parties and hardware devices to each other. Each insurance provider picks one of several implementations of ePA's server-side infrastructure to run for its clients. Currently, there are two approved implementations of this server-side infrastructure.

With the current version of the specificatoin, the overall architecture of ePA heavily relies on Trusted Execution Environments (TEEs). Data processing on the server side is done in plaintext inside TEEs, with some cryptographic key management delegated to a Hardware Security Module (HSM). While attacks on the TEEs are considered in the system, the HSMs are assumed to be perfectly secure, and the system does not include mitigations for a compromised HSM. The primary motivation for plaintext processing seems to be to enable large-scale data analysis for research purposes without requiring consent or cooperation of the people whose records are being processed [W82].

The primary services offered by the server side are authentication services, key escrow, and a database storing the encrypted records themselves. Records are symmetrically encrypted with keys that are derived from system-wide secrets inside an HSM. The primary motivation behind the use of a key escrow service seems to be to enable the creation of a duplicate user ID smartcard in case an enrolled person looses theirs. While the current version of the standard is unclear on the exact mechanism of key derivation, in previous versions of the standard, the escrow service's root

key, a random salt, and the healthcare ID number of the enrolled person was used in SHA256-HKDF. The specification requires that a new root key is generated once a year, but as far as we can tell, record key rollover is not done automatically but is only meant to be done when the *user* requests it, and old root keys must be retained forever to ensure old records can be accessed. Through this lack of automatic key rollover combined with the need to retain root keys indefinitely, attack surface is maximized and incremental compromises of the system over long time spans become possible.

1.1 Previous Analyses

gematik, the state-owned company specifying the system, commissioned several security assessments of the system relating to the key escrow service. Fischlin [69] focuses on the cryptographic dimension of the key escrow service used in an older version of the standard, and is now obsolete. Slany [W235] approaches the system at a higher level, and focuses on the cryptography of the inner protocol layers spoken between the system's components. Industry research organization Fraunhofer SIT was comissioned for a structured, theoretical assessment of attack paths to the system [W72]. We are not currently aware of independent academic security research on the system.

The design and operation of the system have been independently described in detail by civil society activists, who have demonstrated several successful attacks on the system. Tschirsich, Brodowski, and Zilch [W256] demonstrated how they could trivially acquire each of the smartcards as well as the Konnektor necessary for accessing the system. Tschirsich and Kastl [W257] summarize the history of attacks demonstrated on the system and show multiple practical attacks on various parts of the system's implementation.

2 Concerning Cryptographic Engineering Choices

We wish to highlight some of the design choices in the system that we believe stray from current best practice. This is by no means an exhaustive list, and is only meant to underscore why we believe the system deserves more scrutiny.

2.1 Use of Key Escrow

Key escrow describes a concept that was originally devised during the 1990ies out of a fear that the widespread availability of strong encryption would stifle the ability of law enforcement agencies to wiretap communications in the prosecution of crime. At the core of the concept rests the idea that a trusted *key escrow* service should hold a copy of every private key in use. In case the government wants to access one of these keys, the key escrow service can provide this accessAnderson [14] and Jarvis [127].

While key escrow services have been a topic of political debate in decades past, in the cryptographic community, consensus generally is that they are a bad idea since they pose a centralized target for attack, and increase attack surface [4, 5, 14, 220].

To do
Feedback from HS3 reviewer: I feel that this section is a mix-up of critique on the cryptographic design and the approach to privacy protection and data minimisation. How are they linked? I'm missing some discussion here.

Our first concern is the system's general approach of using a key escrow service instead of securely storing the keys inside the system's already existing smart card infrastructure. Like any other key escrow system, this key escrow service poses a centralized security risk. The system's designers made this decision since it was considered important that when an encrypted record must be restored after an insurance ID card is lost, it can be re-created without the cooperation of the healthcare providers holding the primary copies of the person's medical records.

2.2 Cryptographic Design

The system's overall cryptographic design is intentionally kept simple. The standard explicitly mentions that symmetric primitives have been preferred over asymmetric primitives in the core key escrow functions due to the risk of an attack on asymmetric primitives in the long term. Notably, other advanced cryptographic techniques such as secret sharing schemes, oblivious pseudo-random functions, or multiparty computation that could help with the security and privacy of the key escrow service by reducing trust placed in any single component of the service are also absent while the system relies extensively on the engineering-based security guarantees of TEEs and HSMs. Given that the ePA system trusts its HSMs as unconditionally secure, it is unclear what purpose the manual yearly root key renewal serves, especially absent an automatic way to roll over the wrapped record keys.

A consequence of the systems' simple cryptographic design is that the system trusts its components to a large degree. For instance, the system leaks a person's insurance ID number to the key escrow HSM every time

record keys are requested. Along with the timing and frequency of these requests, this leaks information on the person's condition to the key escrow service in an identifiable way.

2.3 A Realistic Attacker Model

We observe that the system as a whole does not appear to be designed to defend against well-resourced adversaries. A series of demonstrated practical attacks on the system, none of which required advanced capabilities, confirm this impression. In Tschirsich and Kastl [W257] summarize a series of successful attacks. Attacks include social engineering resulting in access to copies of smartcards enabling accessing patient records, using misconfigured Konnektor VPN appliances with their local network DMZ and authentication interface exposed on the public internet, circumventing video-based authentication processes resulting in duplicate file keys being provided, classis SQL injection on a backend service maintaining an authentication database, accessing all national patient records through brute-force enumeration of weak identifiers, and several more.

We believe that a system like this must be designed to withstand well-resourced adversaries such as foreign secret services, since the medical data stored in such as information on chronic illness, sexually transmittable disease or severe food allergies has intelligence value. Repeated breaches of national digital infrastructure such as the 2015 breach of the US Office of Personnel Management [W20] or the 2024 compromise of US telecommunications wiretapping systems [W170] demonstrate that such state-sponsored attacks on national digital infrastructure are a realistic concern. A possible scenario in the ePA system would be an foreign secret service gaining access to one of the HSMs storing the systems' root secrets, extracting the root secret by an advanced physical attack, then being able to decrypt captured encrypted health records at will. Similarly, a nation-state adversary might have access to an exploit allowing the compromise of the system's TEEs, which would enable the extraction of any patient records being processed in plaintext inside these TEEs.

2.4 Physical Security

Physical security has received some consideration in the system's specification. First, smart cards are used extensively for authentication. Second, Hardware Security Modules are used in key locations of the system to pro-

cess some cryptographic secrets. The core of the system's key escrow service is implemented inside an HSM that is part of a redundant HSM cluster. However, it is notable that the actual security level required for this HSM is only FIPS 140-2 level 3 [1]. FIPS 140-2 is a US government standard that used to be popular for the specification of HSMs. However, not only has FIPS 140-2 been made obsolete by FIPS 140-3 in 2019 [2], its security level 3 mostly provides logical separation of cryptographic functions from other logic and is not very meaningful in the context of physical attacks. The only physical requirement of FIPS 140-2 level 3 is that the HSM has a hard, opaque coating. This coating is specified to be tamper-evident, but notably no active tamper detection or response features are required by this standard [14]. In contrast to the newer FIPS 140-3 standard and the related ISO/IEC 19790 [W125] as well as ISO/IEC 24759 [W126] standards, FIPS 140-2 does not make any particular requirements regarding resistance to side-channel attacks. The lack of tamper response, unspecified resistance to side-channel attacks and the fact that the ePA specification only requires the long-lived key escrow root key inside the HSM to have 256 bits of entropy lead to an unsatisfactory overall constellation.

3 Conclusion

In conclusion, we observe that in Germany's ePA national medical record database, despite the decade-long standardization and implementation process, several cryptographic compromises ended up in the system's final deployment. Even assuming that nation-scale key escrow is a good idea, the implementation of this key escrow system seems to stray from current best practice. The system uses a secret key with only 256 bits of entropy to derive highly sensitive secret keys for potentially tens of millions of people sharing an insurance provider. The cryptographic design of this escrow system is unsophisticated, ignoring the past three decades in cryptographic developments particularly in multiparty computation (MPC) and other secret sharing techniques in favor of an engineering approach. In the engineering dimension, the system's physical security is only held to the basic level 3 of the obsolete FIPS 140-2 standard, which is considerably less secure than an average credit card payment terminal. The system's root keys are only protected by a "hard, opaque potting material" and no tamper detection and response is required. We estimate that the system poses an attractive

and soft target to nation-state adversaries. The system's shortcomings are made more severe by the fact that the system disproportionally affects the lives of people with low income.

From an academic perspective, it is interesting to see how the ePA ended up in its current state, and the gaps in cryptographic solutions left by academic research that contributed. A fundamental truth in cryptographic engineering is that in the absence of technical checks, political promises are no guarantees of restraint. As such, the degree of trust the ePA system places on organizational measures leads to a concerning overall picture. In particular, the system's extensive reliance on not just conventional HSMs built to long obsolete security standards but also on trusted execution environments that have been broken multiple times highlights the need for new approaches to hardware security that better accomodate real-world use cases.

We believe that Inertial HSMs can address this use case by cleanly separating the physical security primitive into a retargetable design that can be applied to entire servers if needed, and augment or replace technology like conventional HSMs or trusted execution environments to provide high-level hardware security. Before introducing IHSMs in Chapter 4, in the following chapter, we will first complement this chapter's outlook on the state of the art in hardware security with a survey of tamper sensing meshes in a wide range of real world devices.

Web sources

[W10] Amazon. *AWS CloudHSM*. Amazon Web Services, Inc. URL: https://aws.amazon.com/cloudhsm/ (visited on 2025-11-21) (cit. on p. 18).

[W20] Devlin Barrett, Danny Yadron, and Damian Paletta. *U.S. Suspects Hackers in China Breached About 4 Million People's Records, Officials Say*. Wall Street Journal. 2015-06-04. URL: http://www.wsj.com/articles/u-s-suspects-hackers-in-china-behind-government-data-breach-sources-say-1433451888 (visited on 2025-05-15) (cit. on p. 23).

[W72] Fraunhofer SIT. *Abschlussbericht Sicherheitsanalyse Des Gesamtsystems ePA Für Alle*. 2024-08-09. URL: https://www.sit.fraunhofer.de/fileadmin/dokumente/studien_und_technical_reports/

Abschlussbericht_Sicherheitsanalyse_ePA_fuer_alle_Fraunhofer_SIT.pdf (visited on 2025-05-16) (cit. on p. 21).

[W82] Gematik. *Whitepaper Datenschutz und Informationssicherheit in der Telematikinfrastruktur*. 2025-07. URL: https://www.gematik.de/media/gematik/Medien/Newsroom/Publikationen/Informationsmaterialien/gematik_Whitepaper_Datenschutz_web_20250707.pdf (visited on 2025-11-21) (cit. on p. 20).

[W83] gematik. *Spezifikation Aktensystem ePA für alle v1.4.1*. 2025-05-09. URL: https://gemspec.gematik.de/docs/gemSpec/gemSpec_Aktensystem_ePAfueralle/latest/ (visited on 2025-05-16) (cit. on p. 19).

[W84] gematik. *Spezifikation Schlüsselgenerierungsdienst ePA v1.6.0*. 2023-03-31. URL: https://gemspec.gematik.de/downloads/gemSpec/gemSpec_SGD_ePA/gemSpec_SGD_ePA_V1.6.0.pdf (visited on 2025-05-26) (cit. on p. 19).

[W85] gematik. *Übergreifende Spezifikation Verwendung Kryptographischer Algorithmen in Der Telematikinfrastruktur v2.28.1*. 2024-02-23. URL: https://gemspec.gematik.de/downloads/gemSpec/gemSpec_Krypt/gemSpec_Krypt_V2.28.1.html (visited on 2025-05-16) (cit. on p. 19).

[W86] gematik. *Übergreifende Spezifikation Verwendung kryptographischer Algorithmen in der Telematikinfrastruktur v2.40.0*. 2025-03-28. URL: https://gemspec.gematik.de/downloads/gemSpec/gemSpec_Krypt/gemSpec_Krypt_V2.40.0.pdf (cit. on p. 19).

[W90] *Github Repository: eRP-FD/Vau-Hsm*. URL: https://github.com/eRP-FD/vau-hsm/tree/master (visited on 2025-05-16) (cit. on p. 19).

[W92] Google. *Cloud HSM | Cloud Key Management Service*. Google Cloud Documentation. 2025-11-13. URL: https://docs.cloud.google.com/kms/docs/hsm (visited on 2025-11-21) (cit. on p. 18).

[W115] IBM. *Cloud HSM*. 2016-05-01. URL: https://cloud.ibm.com/catalog/infrastructure/cloud.ibm.com/catalog/infrastructure/hardware-security-module (visited on 2025-11-21) (cit. on p. 18).

[W125] *ISO/IEC 19790:2025*. ISO. URL: https://www.iso.org/standard/82423.html (visited on 2025-05-15) (cit. on pp. 5, 24, 38).

[W126] *ISO/IEC 24759:2025*. ISO. URL: https://www.iso.org/standard/82424.html (visited on 2025-04-08) (cit. on pp. 5, 24, 115, 119).

[W138] Marie-Claire Koch. *More and More Experts Warn against Electronic Patient Records*. heise online. 2025-01-10. URL: https://www.heise.de/en/news/More-and-more-experts-warn-against-electronic-patient-records-10235907.html (visited on 2025-05-26) (cit. on p. 18).

[W139] Marie-Claire Koch. *Noch viele Unklarheiten bei der elektronischen Patientenakte*. heise online. 2025-05-08. URL: https://www.heise.de/hintergrund/Elektronische-Patientenakte-Welche-Unklarheiten-es-noch-gibt-10377344.html (visited on 2025-11-28) (cit. on p. 18).

[W170] Joseph Menn. *Chinese Government Hackers Penetrate U.S. Internet Providers to Spy*. The Washington Post. 2024-08-27. URL: https://www.washingtonpost.com/technology/2024/08/27/chinese-government-hackers-penetrate-us-internet-providers-spy/ (visited on 2025-05-15) (cit. on p. 23).

[W172] Microsoft. *Overview of Azure Cloud HSM*. URL: https://learn.microsoft.com/en-us/azure/cloud-hsm/overview (visited on 2025-11-21) (cit. on p. 18).

[W213] *Quote Origin: The Most Dangerous Phrase Is: "We've Always Done It That Way" – Quote Investigator*. 2014-11-27. URL: https://quoteinvestigator.com/2014/11/27/always-done/ (visited on 2025-10-22) (cit. on p. 17).

[W235] Wolfgang Slany. *Sicherheitsanalyse zur Sicherheit der kritischen Komponenten der elektronischen Patientenakte nach §291a SGB V*. 2020-03. URL: https://www.gematik.de/media/gematik/Medien/Newsroom/Presse/Dokumente/Sicherheitsanalyse_TU_Graz_zur_ePA_mit_Vorwort_der_gematik.pdf (visited on 2025-05-15) (cit. on p. 21).

[W248] Thales. *Luna Network Hardware Security Modules*. URL: https://cpl.thalesgroup.com/encryption/hardware-security-modules/network-hsms (visited on 2025-11-21) (cit. on p. 18).

[W256] Martin Tschirsich, Dr med Christian Brodowski, and Dr André Zilch. *"Hacker Hin Oder Her": Die Elektronische Patientenakte Kommt!* 2019-12-27. URL: https://media.ccc.de/v/36c3-10595-hacker_hin_oder_her_die_elektronische_patientenakte_kommt (visited on 2025-05-15) (cit. on p. 21).

[W257] Martin Tschirsich and Bianca Kastl. *„Konnte Bisher Noch Nie Gehackt Werden": Die Elektronische Patientenakte Kommt - Jetzt Für Alle!* 2024-12-27. URL: https://media.ccc.de/v/38c3-konnte-bisher-noch-nie-gehackt-werden-die-elektronische-patientenakte-kommt-jetzt-fr-alle (visited on 2025-05-15) (cit. on pp. 21, 23).

[W260] Utimaco. *Use Cases.* URL: https://utimaco.com/use-cases (visited on 2025-11-21) (cit. on p. 18).

[W261] Utimaco. *What Is Cloud HSM?* 2025-09-10. URL: https://utimaco.com/service/knowledge-base/hardware-security-modules/what-cloud-hsm (visited on 2025-11-21) (cit. on p. 18).

[W273] *What Is a Cloud HSM?* URL: https://www.entrust.com/resources/learn/what-is-cloud-hsm (visited on 2025-11-21) (cit. on p. 18).

[W277] *WikiQuote: Grace Hopper.* 2025-04-08. URL: https://en.wikiquote.org/wiki/Grace_Hopper (visited on 2025-10-22) (cit. on p. 17).

References

[1] (US) National Institute of Standards and Technology. *Security Requirements for Cryptographic Modules.* Federal Information Processing Standard (FIPS) 140-2. U.S. Department of Commerce, 2002-12-03. DOI: 10.6028/NIST.FIPS.140-2 (cit. on pp. 2, 4, 24, 115, 119).

[2] (US) National Institute of Standards and Technology. *Security Requirements for Cryptographic Modules.* Federal Information Processing Standard (FIPS) 140-3. U.S. Department of Commerce, 2019-03-22. DOI: 10.6028/NIST.FIPS.140-3 (cit. on pp. 2, 4, 24, 38, 66).

[4] Hal Abelson et al. "The Risks of Key Recovery, Key Escrow, and Trusted Third-Party Encryption". In: *World Wide Web J.* 2.3 (1997-06-01), pp. 241–257. ISSN: 1085-2301 (cit. on pp. 2, 22).

[5] Harold Abelson et al. “Keys under Doormats”. In: *Commun. ACM* 58.10 (2015-09-28), pp. 24–26. DOI: 10.1145/2814825 (cit. on pp. 2, 22).

[14] Ross Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 1st ed. Wiley, 2020-12-22. DOI: 10.1002/9781119644682 (cit. on pp. 2, 22, 24, 36–38, 43, 64, 75, 77, 78, 92, 118).

[69] Marc Fischlin. *Kryptographische Analyse Spezifikation Schlüsselgenerierungsdienst ePA*. Technische Universität Darmstadt, 2021-12 (cit. on p. 21).

[127] Craig Jarvis. *Crypto Wars: The Fight for Privacy in the Digital Age: A Political History of Digital Encryption*. 1st ed. CRC Press, 2020-12-14. ISBN: 978-1-00-312367-5 (cit. on pp. 2, 22).

[220] Phillip Rogaway. “The Moral Character of Cryptographic Work”. In: *Advances in Cryptology*. ASIACRYPT 2015. Vol. 9452 & 9453. LNCS. Auckland, New Zealand: Springer, 2015, p. XVIII. DOI: 10.1007/978-3-662-48800-3 (cit. on pp. 2, 22, 217).

Chapter 3

Active Tamper Sensing in the Wild

> Bypassing a PAL [atomic bomb ignition code lock] should be about as complex as performing a tonsillectomy while entering the patient from the wrong end.
>
> – *An unnamed atomic bomb designer [26]*

Contents

Inertial Hardware Security Modules are the latest link in a series of developments bringing hardware security primitives from niche military cipher machines to mass-market applications. The tamper sensing technology that forms the primary line of defense in such physical security systems goes back more than a century, with the earliest tamper sensing meshes being used in the late 19$^{\text{th}}$ century, around the widespread commercialization of electricity. Today, active tamper sensing meshes are used in a wide array of devices ranging from card payment terminals to atomic bombs.

In this chapter, we will start with a brief history of tamper sensing meshes. Complementing our historical analysis, we will present the results of a survey of a range of real-world devices that use tamper sensing meshes and we will examine their implementation. We will analyze the gaps left by the current state of the art in commercial practice, and evaluate how Inertial HSMs could close these gaps to make secure hardware accessible to a wider range of applications. The contributions in this chapter are as follows:

- We provide a historical overview of uses of tamper sensing meshes.
- We provide the first large-scale analysis of real devices incorporating tamper sensing meshes in the academic record.
- We create a taxonomy of practical construction techniques and provide both detailed analyis and photos illustrating them.
- From our sample, we extract several design patterns that can be applied to increase the security of a design.
- We note security flaws in several of our samples.
- We provide the results of Computed Tomography (CT) imaging of multiple samples, and we evaluate their impact on tamper sensing mesh security.

1 The History of Tamper Sensing Meshes

Tamper sensing meshes offer many degrees of freedom in their design ranging from the precise conductor layout, through the manufacturing technology of the mesh and how it is wrapped around the payload during manufacturing up to their monitoring circuitry. As a result, manufacturers across

application domains from datacenter appliance HSMs to card payment terminals have historically used patents on parts of their tamper sensing mesh implementations as a means to prevent copying of their designs [P215, P103, P47, P104, P204]. The basic principle of modern tamper sensing meshes is to reliably detect physical intrusion using an embedded looped conductor to cover a surface. This concept traces back at least as far as 1870 [P121, P120], when it was applied to the protection of bank vaults from robbers attempting to dig, drill and saw through the vault's floor and walls. Even multi-layer, orthogonal tamper sensing meshes are documented as far back as 1902 [P244]. Using printed circuits instead of wires for this purpose occurs in literature as soon as printed circuit technology finds widespread commercial adoption in the 1960ies [P101]. The history of more HSM-like devices begins in the 1990ies with the widespread adoption of cryptography in commercial applications [P136, P129, P59, P41, P42, P61, P49, P33, P48, P207] when instead of protecting an entire device it became feasible to create a protected cryptographic coprocessor.

1.1 Use by the US Military

One early practical uses of tamper sensing meshes for information security as opposed to the security of some physical good is documented in notes on a series of lectures given by Dr. David G. Boak, a specialist in communications security and signal intelligence at the US National Security Agency [29, 28]. In this lecture series, Boak mentions that around World War II, the US became concerned about the security of their ciphering machines, which at the time were large, fridge-sized electro-mechanical contraptions. Initially, simple safes were used to protect those devices—however, as Boak notes, the US was well aware that they could not build a safe that a well-equipped specialist could not break open within an hour. As a solution, the NSA started development on what we would today call a Hardware Security Module by encapsulating a crypto coprocessor in a tamper sensing envelope. Boak observes that as a tamper response, reliably zeroizing the cryptographic keys would be sufficient. Today, this approach is universally taken. Boak does note several other ways to penalize an intrusion attempt, including raising a remote alarm or–even more exciting–exploding the device.

1.2 Use in Nuclear Weapons

Communications security was not the earliest use of tamper sensing membranes in the US military, with Boak mentioning HSMs still being under development in the second volume of the lecture series, dated 1972. An earlier reference to such systems can be found in literature on Permissive Action Links (PALs) for nuclear weapons. In US military terminology, a PAL is a chain of locked, tamper-proof systems required to trigger the detonation of a nuclear weapon. PALs were developed as a consequence of nuclear weapons being stationed in countries allied with the US during the cold war. The concern was that the host country might forcibly assume control over the US nuclear weapons stationed on their soil. The stated goal of PALs is to protect the weapon from use without a secret passcode known only to US military command. To achieve this goal, PALs will lock themselves when incorrect codes are entered. To protect against both intentional tampering aiming to circumvent the PAL, as well as against accidential detonation under extreme environmental conditions, PALs are designed such that any tampering attempt as well as any environmental deviation will be sensed by the PAL, and will lead to the weapon being destroyed in a less harmful way that does not cause the full-scale nuclear explosion that the weapon is capable of. This goal is achievable in practice since nuclear weapons are reportedly very sensitive to the timing of their primary explosive charges, as the nuclear payload only produces a full-scale detonation when triggered in just the right way.

While it is difficult to date, Carter et al. [39] specifically mention a tamper sensing membrane being used in US PALs. Given the nature of the matter, it is safe to assume that this technology will have been in use for some years at the point it was being discussed in an unclassified, civilian book on nuclear armament control.

1.3 Use in Nuclear Safeguards

Besides being used in nuclear weapons, tamper sensing systems have another, more peaceful application in the nuclear field. In 1957, the International Atomic Energy Agency (IAEA) was founded to coordinate and verify that civilian nuclear energy installations are not used for military purposes. A core part of the IAEA's tasks is observing the operations at civilian nuclear installations through inspections and through a variety of permanently deployed sensors to track the history of nuclear material passing through

these facilities.

When using sensors to monitor treaty compliance, the IAEA has to consider the possibility of a host state tampering with its sensors to abuse nuclear material without being noticed. Historically, the IAEA has responded to this threat by the extensive use of tamper-indicating enclosures and of seals[1]. In both systems, the approach taken is that the enclosure or seal is treated similarly to what these days, in computing we call a Physical Unclonable Function (PUF). The concept of a PUF centers on electronic component manufactured such that random manufacturing variations can later be measured by the finished circuit. The core idea is that since these manufacturing variations are random, they can be used as a source for cryptographic entropy. Furthermore, the concept is based on the assumption that these manufacturing variations cannot be controlled, hence making the device *unclonable*.

Similar to a PUF, in the IAEA's application an enclosure or seal is manufactured in a process that leaves an unpredictable and uncontrollable pattern of manufacturing variations such as surface imperfections. A process used in the IAEA is to package devices in aluminium enclosures passivated in a bright color, which leaves a random, microscopic pattern of pits in the surface from the etching step. Before such a device is deployed in the field, it is precisely measured from all sides. Later on, after field deployment, its integrity can then be checked by comparing its current state to these initial measurements. The underlying assumption is that drilling or cutting into something like a metal enclosure will leave detectable traces, and that perfectly replicating an object including features such as minute surface imperfections is infeasible even to a nation state [122].

With smarter electronics becoming more affordable in both monetary and in power budget, over the decades, other active tamper sensors have received attention as well. The IAEA reports on attempts at burying sensors such as piezoelectric transducers or optical fibers inside an enclosure's walls to detect tampering, but states that these efforts have not yielded practical

[1]Note that in IAEA terminology, both tamper detection and tamper evidence are combined into the term "tamper indication". The IAEA distinguishes between active tamper indication, which we conventionally call tamper detection, and passive tamper indication, which we conventionally call tamper evidence. Tamper indicating devices include seals, but also the aforementioned uniquely characterizable enclosures, which IAEA terminology calls intrinsically tamper-indicating. An example for an active tamper indicating device would be a seismic sensor at the bottom of a borehole that has been back-filled with concrete such that any attempt to reach the sensor would be well-visible in the sensor's own readings [234].

results primarily due to cost concerns. In contrast to these sensors, the IAEA's Electro-Optic Sealing System (EOSS) uses a flexible tamper sensing mesh that contains some sort of conductive traces in the same way it is used in contemporary hardware security modules to detect attempts at drilling or cutting into the system [122, 252]. Unfortunately, no information on the precise construction of the tamper sensing mesh such as materials used or structure sizes are publically available.

1.4 Commercial Use

Commercially, tamper sensing meshes have entered widespread use beginning around the turn of the millennium, initially in then-new HSMs, cryptographic coprocessors primarily aimed at the financial industry [14]. Today, their use in finance has spread from HSMs in datacenters and ATMs to the ATM pin pads themselves, which encrypt the customer's PIN right at the source, as well as in all kinds of card payment terminals.

HSMs are used for highly sensitive operations even outside of the financial industry, although their adoption is hampered by their high cost. In this chapter, we will analyze a commercial HSM that was used in the key management infrastructure of a premium TV provider as one example of such uses. Examples of other applications include mail franking machines, where they are used to protect the credit counter and franking data, with one such unit analyzed in this chapter. Furthermore, we have identified several models of key safes that in Germany are mounted externally on public buildings to provide keys to emergency services, and which include tamper sensing meshes on their door and interior walls to detect attempts at drilling into them [W232, W144]. Finally, we have found a processing unit used in a series of mid-2000s era slot machines in Germany that includes a tamper sensing mesh, presumably to prevent modification or cloning. This device will also be analyzed later in this chapter.

2 Tamper Sensing Mesh Design Principles

The manufacturing technology of a tamper sensing mesh is a critical factor in its security. While in many applications, meshes manufactured from off-the-shelf processes such as Flexible Printed Circuit (FPC) processes are used, these processes tend to be optimzed to maximize the robustness of the produced circuits to mechanical stress. In contrast, the ideal tamper

sensing mesh is exactly as robust as it needs to be not to be destroyed accidentially during normal handling, but should not be more robust than that. As a result, more secure meshes tend to be manufactured in bespoke manufacturing processes [116, 117, P121, P105, P196, 264, 236].

One more widely cited tamper sensing mesh implementation is a commercial product developed by IBM in collaboration with chemical company W. L. Gore & Asscociates Inc. This product is used in IBM's datacenter HSM products up to approximately 2020 [195, 14, 236]. It uses a stack of multiple layers of a clear, flexible plastic substrate on which carbon-based traces are printed. Vias, i.e. contacts between layers, are made by laser cutting small holes into the substrate before the traces are printed. The flexible circuit layers are joined with a opaque black, stretchy glue and are embedded in an elastic opaque resin after installation. The plastic substrate foil is thinner and significantly less resistant to tearing than plastic substrates commonly used in the electronics industry for applications like key pads and circuit boards, which improves its security against tampering. It is clear that both the glue fusing the foil layers together and the resin that the mesh is embedded inside are co-designed with the carbon trace material such that the trace material adheres well to both, leading to the traces being destroyed when either are peeled off.

The design of these IBM/Gore meshes is documented in an extensive list of patents, mostly under IBM's name. Its basic construction and layout has not changed much since the early 1990ies [P161, P160].

Concluding this brief history of tamper sensing meshes, we find that they were initially developed for sensitive military applications, and their use in civil applications is a recent phenomenon. The implementation of tamper sensing meshes in civil applications was likely catalyzed by two advancements in electronics. First, electronic components became less expensive and more integrated reducing the cost overhead of tamper sensing circuits. Second, the mass-scale adoption of PCB and FPC production processes enabled their use as inexpensive, high-resolution substrates for such meshes.

2.1 Monitoring Circuit Approaches

Tamper sensing meshes are most effective when they are continuously monitored using a backup power supply while the rest of the system is powered off. In practice, the main challenge with continuous monitoring of tamper

sensing meshes is in the design of the monitoring circuit. A large portion of industry attention has been spent on designing low-power monitoring circuits that are sensitive to tampering with the mesh while using little enough power to enable years of operation from a battery. Commonly, one or two cylindrical or large coin cell Lithium primary batteries are used, providing in the order of 10 W h over their lifetime[109]. Broken down to an unpowered storage life of e.g. 5 years, this corresponds to a maximum average power consumption of less than 230 µW.

To achieve low power consumption, a popular technique known since at least 1902 [P244] and still used today [P41, P215] is to measure the deviation of the mesh's end-to-end ohmic resistance from its baseline value. This measurement can be implemented either by directly comparing a mesh trace's resistance with a reference resistor, or using a Wheatstone bridge. Bridge circuits were already used in early tamper sensing mesh implementations [P62, P101, P54] since they make it possible to detect small changes in the mesh's resistance with little complexity.

2.2 Other Tamper Sensing Techniques

Besides tamper sensing meshes, environmental sensors such as temperature or light sensors are frequently used as a secondary line of defence in HSMs and similar devices. By placing such sensors in the device and verifying the device is within its nominal operating environment, tampering can be made less convenient. Modern security standards often mandate the implementation of at least a temperature sensor to prevent cold-boot attacks on a device [2, W125]. A multitude of other sensors have been proposed, including vibration sensors, light sensors, magnetometers, and radiation sensors such as X-ray sensors have been proposed. While the implementation cost of most sensor types is low, each additional environmental sensor comes with an increased false alarm rate [14].

3 A Survey of Meshes in the Wild

In this section, we will examine a large sample of recent devices that include tamper sensing meshes to gain an understanding of how they are implemented, and what security level they are targeted towards. Since we were unable to acquire a nuclear weapon for our research, we limited our survey to commercial devices. While we analyzed devices across a broad spectrum

of applications, our survey includes a large variety of card payment terminals, which represent the most varied class of device incorporating such meshes.

3.1 Specimen Selection

Given their niche applications and high cost, devices incorporating tamper sensing meshes tend to be hard to find. For this survey, we chose 30 total devices including 23 different models of card payment terminals, and 7 other devices. Some devices were procured by intercepting electronic waste, while most were sourced from ebay in Februrary and March 2025. The majority of these were sold by electronic waste recycling companies. A complete list of our specimens can be found in Table 3.1. External photos of each device are shown in Figure 3.1 and internal photos are shown in Figure 3.2. In the following sections, we will go into detail on the classes of devices we selected for this study.

Card Payment Terminals

Card payment terminals commonly include advanced tamper sensing features to discourage physical attacks such as skimming that aim to exfiltrate card data and PINs entered by the customer. The Payment Card Industry Security Standards Council (PCI SSC), an association of all major western credit card network operators assumes the role of the de-facto standardization organization in the card payment space. Due to the international scale of the large credit card networks, almost all payment terminals on the market irrespective of their country of origin are certified under PCI SSC standards. Adding on to PCI's ecosystem impact, its security standards are thought out well.

One reason for the high level of physical security standards in card payment applications both on the client side (payment terminals) and on the server side (HSM appliances) is that the finance industry has been reluctant to adopt modern cryptography. Not only are modern cryptographic protocols like secure Multiparty Computation (MPC) or Zero-Knowledge Proofs (ZKPs) not commonly used. Even asymmetric cryptography has only been adopted reluctantly, and ancient ciphers such as Triple DES are still commonly referenced in industry standards [202]. As a result, increased hardware security is necessary to safeguard weak symmetric keys, compensating for the systems' modest cryptographic security.

ID	Device	Manufacturer	Type code	Year
H01	PED	Verifone	VX 570	ca. 2010
H02	Slot machine CPU module	Merkur / ADP Gauselmann	Sam 12 EC2	ca. 2012
H03	EPP	Sagem	USA1315-4240 R1A	2014
H04	EPP	Sagem	USA1316-5120 R1A	2007
H05	PED	Xac	xAPT-103	2014
H06	PED	Ingenico	iCT250-11T1860A	2016-17
H08	PED	Sagem	NOR4100-4220 R1A	2012
H09	PED	Hypercom	M4230	2010
H10	PED	Worldline	YOMANI XR	2016
H11	PED	Banksys	C-ZAM Smash Portable	2004
H12	PED	Hypercom	Optimum P2100	2010
H13	PED	Ingenico	iCT 220-11T2938A	2016
H14	PED	Verifone	H5000	2016
H15	PED	Verifone	MX 925	2018
H16	PED	Verifone	V200c CTLS	2021
H17	PED	Verifone	VX 680	2014
H18	PED	Ingenico	i7910	2010
H19	PED	Banksys	XENTA	2004-2011
H20	PED	Verifone	VX 520 3G	2017
H21	PED	Verifone	V400m Plus 4G	2018
H22	PED	Ingenico	Move 3500	2020
H23	PED	Ingenico	iPP 350-11T1718A	2015
H24	PED	Ingenico	iWL255-01T2117A	2016
H25	Franking Machine	Neopost	IJ-25	ca. 2001
H27	PED	Sumup	AIR1E205	2021
H28	EPP	NCR	5814 UEPP	2019
H29	HSM	SafeNet	VBD-05	2018
H30	HSM	Irdeto	Mayflower-IDX/C201	2011
H31	PED	SumUp	SumUp 3G	2019
H32	PED	SumUp	SumUp Air	2022

Table 3.1: The specimens we dissected in our survey. PED stands for *Pin Entry Device*, the industry term for card payment terminals that have sufficient security to handle credit card PINs. EPP stands for *Encrypting Pin Pad*, the type of keypad used for pin entry on ATMs. HSM stands for Hardware Security Module.

Figure 3.1: External photos of all survey specimens.

Since card payment terminals are widely deployed, many different models from various manufacturers are available. Each manufacturer tends to have their own, patented tamper sensing implementation. Being manufactured at scale, card payment terminals are cost-sensitive devices, which is reflected in the construction of their tamper sensing implementations.

HSM Appliances

When credit card payments are handled on the web as opposed to in a physical store, HSMs are used in data centers to handle plaintext payment data such as credit card numbers. Such HSM appliances are usually standalone rackmount devices and are used across application domains. Depending on the application, these HSMs can be programmed with custom code, or can be used as coprocessors through an API [W249]. In practice, the standalone appliances are just low-end computers in a rackmount enclosure that expose the API of an internal HSM add-in card to the network. In this survey, we obtained two devices labelled as HSMs. We were only able to procure two such devices since they are expensive, and we found that even used specimens of older models are usually listed for several hundreds to several thousands of Euro. Unfortunately, one of the devices we obtained did not contain any security meshes in its case, and thus would not provide adequate protection against advanced attacks. The other specimen we procured was a 2011 model Utimaco CryptoServer LAN. Our unit was a white-label variant procured by premium TV encryption technology provider Irdeto, presumably used in Germany to produce cryptographic key streams for TV signal encryption. We bought the device from a recycling company specialized on datacenter components. The device was sold with any HDDs removed. The device consisted of an older mainboard for embedded applications containing an Intel Core 2 Duo-brand processor and 2 GiB of DDR2 RAM, which was connected to the HSM add-in card through PCI. The device contained a small Lithium backup battery on the add-in card, and another, larger battery in an enclosure at the front of the device that was connected to the card through a cable. The device did not contain any obvious case intrusion sensors.

ATM Encrypting Pin Pads

ATMs are built in a modular construction approach. Physically, the enclosure of an ATM is not its only security barrier. Besides the enclosure,

there are two security barriers worthy of note. First, the bank notes in the machine are stored in an automatic cash dispenser that is built into a traditional vault inside the machine. This vault primarily acts as a mechanical barrier to discourage theft, but it also often includes tamper sensors that activate an Intelligent Banknote Neutralisation System (IBNS) [W19, W65, 198]. The IBNS is designed to spread hard-to-remove ink over the bank notes inside the vault when tampered. The permanently stained bank notes are not accepted by banks or retailers anymore.

Besides the vault, the another security barrier is located inside the ATM's pin pad. While all communication with the customer's card passes through an end-to-end encrypted channel from the bank's backends into the card's smartcard IC, the customer must necessarily enter their pin in plain text. To prevent leakage of the plaintext PIN, the PIN is encrypted inside the PIN pad itself. To this end, the PIN pad contains a microcontroller handling the encryption [14]. Often, both the circuit board containing the PIN pad's keyboard matrix and this microcontroller are shielded by a tamper sensing mesh to prevent physical attacks such as the installation of a skimming device that would record and transmit the plaintex PIN.

We acquired three different EPPs for analysis: Two designed by Sagem and apparently re-sold as a whitelabel product by Cryptera and Diebold, respectively, and one made by and branded NCR. All three devices have robust stainless steel front cases, and are built in a sandwich construction of several layers of steel sheets and PCBs.

Other miscellaneous devices

Sometimes, tamper sensing meshes show up in other types of devices. We acquired two such devices. First, we acquired a Neopost mail franking machine, a type of device that is used to directly print a code on an envelope that replaces a conventional postage stamp. Since in businesses handling large volumes of mail these devices were routinely charged with large sums of money in postage, such devices have security features ranging from physical seals on their enclosure to full security meshes encasing their CPU modules. In case of Neopost, we are aware of one online source showing a security mesh inside one such device [W173], but we found that our older specimen only contained a sturdy cast zinc case that was welded shut with a spring-loaded lid switch inside. The other miscellaneous device we found is a broken CPU module from a German slot machine manufacturer. While

it would be reasonable to assume this type of device might include active tamper sensing features to enforce state gambling regulations, other slot machine manufacturers seem not to use tamper sensing in their systems so the more likely reason is DRM. Our specimen included both a tamper sensing mesh as well as a semiconductor junction light sensor inside of a sealed sheet metal enclosure.

3.2 Methodology

In this survey, we aim to create a comprehensive taxonomy of tamper sensing mesh construction methods across a range of devices. To this purpose, we proceeded by first photographing every test specimen from multiple angles, then disassembling them. After disassembly, we photographed each major component. Figure 3.2 shows a selection of these photos showing the major internal components of the devices. After photos were taken, we proceeded with destructive techniques where necessary to understand the devices' use of tamper-sensing meshes. We took microscope photos where we found interesting small structures. PCBs were sectioned using a sanding drum attachment on a Dremel rotary tool. Potted modules were disassembled using milling, cutting and prying, and applying heat from a heat gun as necessary to soften polymer compounds and to break glue joints.

To do
Actually assemble the supplementary material and include all photos

3.3 Results

In the following sections, we will list some observations we made while dissecting our specimens. A complete set of internal pictures and micrographs of selected components that goes beyond the following description is available in the supplementary material to this thesis.

Mesh materials.

We found meshes constructed from rigid PCBs (e.g. specimens H02, H03 and H08) as well as a number of FPC processes. Tamper sensing meshes constructed from PCBs sometimes used parts of an existing PCB (e.g. specimens H03 and H10), and sometimes additional PCBs only containing a mesh were added (e.g. specimen H02 and H08). In some samples (e.g. specimens H08 and H18), multiple rigid PCB meshes were assembled in a house of cards fashion to enclose a card slot. All flexible meshes that we found with the exception of the Utimaco HSM appliance's HSM card (specimen H30) were clearly manufactured either entirely or mostly in standard

Figure 3.2: Internal overview photos of the survey specimens.

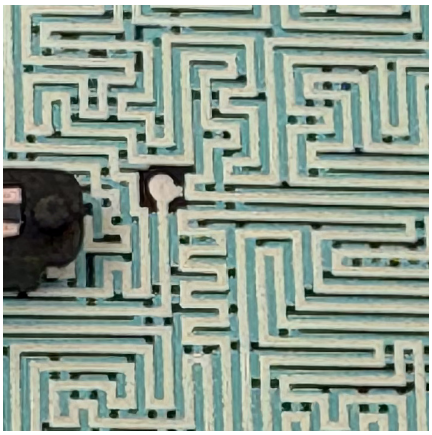

(a) Offset layers for more complete coverage (specimen H12).

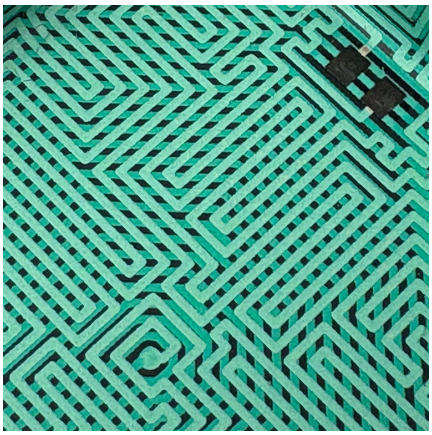

(b) Orthogonal patterns on subsequent layers (specimen H14).

(c) Combining orthogonal layers with area-covering pattern (specimen H30).

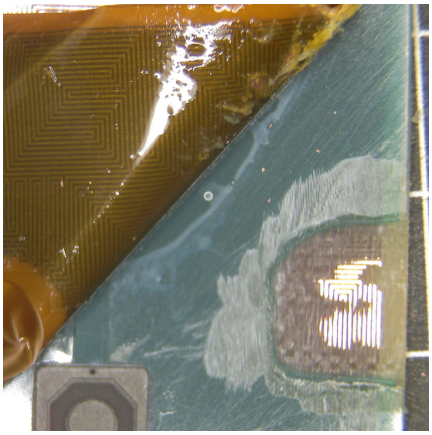

(d) Spacing mesh layers apart to constrict angular freedom of an attack tool (specimen H28).

Figure 3.3: Mesh trace layout approaches for multi-layer meshes.

processes. We found printed silver ink (e.g. specimen H12) and printed carbon ink-based foils (e.g. specimen H09) similar to those used for membrane keyboards, as well as conventional photolithographically etched copper/polyimide FPCs (e.g. specimens H03, H04 and H08). Overall, etched PCBs showed better resolution compared to silkscreen-printed meshes. Feature size for both rigid and flexible etched PCB meshes was generally in the order of 100 µm to 200 µm, while feature size for screen printed foil meshes was coarser at between 500 µm to 3000 µm. In contrast to these standard processes, the Utimaco HSM used a mesh foil that is manufactured in a proprietary, bespoke process by Gore.

Mesh layout.

A key goal in tamper sensing mesh design is to avoid any gaps in coverage. In single-layer meshes, gaps between adjacent mesh traces cannot be avoided, and provide an easy approach for an attack. In multi-layer meshes, these structure size-dependent gaps can be mitigated in multiple ways as shown in Figure 3.3. In the following list, we will address several common structural features that we observed across samples.

1. **Offset patterns.** In a two-sided foil mesh, most of the gaps between adjacent traces can be covered by simply offsetting the pattern by one structure size in both axes between the foil's top and bottom layers as shown in Figure 3.3a. Depending on the mesh layout, only a small number of point-shaped gaps remain at corners in mesh traces on one of the layers. The number of these gaps can be reduced by reducing the number of misaligned corners between both layers for instance by choosing a systematic serpentine or spiral trace layout.

2. **Orthogonal patterns.** In some other specimens, the manufacturer chose the opposite approach of keeping the mesh pattern mostly orthogonal on the mesh's two layers as shown in Figure 3.3b. While this leads to a larger amount of gaps compared to offset patterns as described above, it also reduces the largest gap size to about one structure size by one structure size.

3. **Combined approaches.** Figure 3.3c shows the layout of a Gore tamper sensing mesh foil used in an Utimaco HSM. This mesh consists of two foil layers bonded to each other. The outer foil is patterned on both sides with a sparse pattern of thin serpentine traces with the patterns on both layers being orthogonal to each other. Both patterns are oriented at a 45° angle relative to the sides of the rectangular enclosed volume. The inner foil is only patterned on one side, and contains a thicker serpentine trace laid out in a zigzag pattern. The two foil layers are aligned such that no gaps remain between the layers.

To do
sample number here and below (ingenico)

4. **Using layer spacing.** Figure 3.3d shows how an ATM Encrypting Pin Pad (EPP) implemented the mesh on its keypad. Off-the-shelf metal snap dome contacts were used on the surface of a conventional rigid PCB to create the keys. On top of the rigid PCB and contact

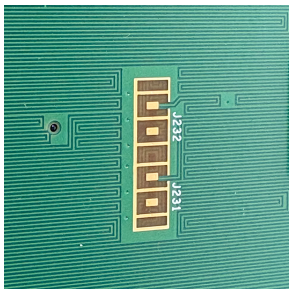

(a) Standard photolithographic copper PCB process on rigid FR-4 fiberglass substrate (specimen H10).

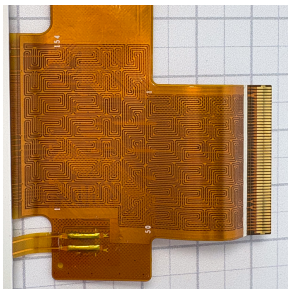

(b) Standard photolithographic copper PCB process on flexible polyimide substrate (specimen H15).

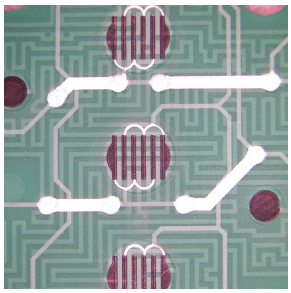

(c) Screen printing process using silver ink with some carbon ink contact pads for embedded buttons (specimen H14).

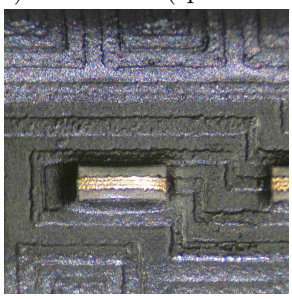

(d) Laser direct structuring using electroless gold plating (specimen H32).

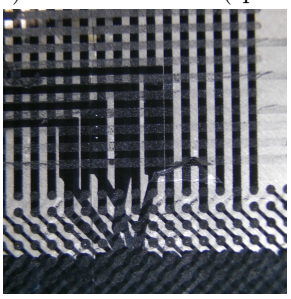

(e) Screen printing process using carbon ink (specimen H30).

Figure 3.4: Materials and manufacturing processes used for mesh traces and contacts.

domes, a two-layer copper/polyimide FPC with an additional polyimide cover layer was glued down. Meshes were placed on both layers of the FPC, as well as on one internal layer of the rigid PCB. The resulting structure had the FPC mesh layers separated from the rigid PCB mesh layer by several hundred micrometers of the rigid PCB's substrate. The meshes on both the FPC and the rigid PCB used a structure size of 150 µm. The vertical separation between the two meshes was several times that structure size, which limits the possible angles an attack tool could be inserted through both mesh layers.

To do
FIXME: Add scale / structure size to photos?

Contact and trace construction.

Regular Printed Circuit Boards are frequently used to implement tamper sensing meshes as shown in Figure 3.4a. PCB production is a highly advanced, large-scale industry and PCBs are inexpensive, commodity products. PCBs can be manufactured with many layers, at almost arbitrary

total thickness, and offer small structure sizes enabling the creation of fine features down to approximately 100 µm even on commodity processes. The primary disadvantage of using PCBs to implement tamper sensing meshes is that PCBs are fundamentally designed to be as robust as possible. The traces on the top of a PCB are etched from a thick (usually 35 µm on the outer layers) copper foil adhered to the PCB substrate. As a result, the PCB and the traces on its surface are easy to manipulate by hand using tools like knives and techniques like soldering. For a tamper sensing mesh, trace patterns manufactured to be more fragile might be advantageous. Additionally, standard PCBs are made using a rigid FR-4 fiberglass/epoxy substrate. Since a tamper sensing mesh must often enclose all sides of a payload, flexible foils offer benefits over rigid PCBs.

Figure 3.4b shows an FPCs produced in a standard commercial process similar to PCB production. In FPCs, a copper foil adhered to a substrate is etched, but the substrate here usually is a thin foil made from polyimide, an orange, temperature-resistant polymer that survives common reflow (hot air) soldering temperatures. In contrast to rigid PCBs, FPCs are usually limited to no more than four layers before losing flexibility. Flexible PCBs are often used for tamper sensing meshes that wrap around a payload, but they come with the same limitation as standard PCBs: Due to their robust substrate and thick copper layers, they are easily manipulated by hand.

Figure 3.4c shows an FPC created in a different process. Here, instead of photolithographically etching a continuous copper foil adhered to a flexible substrate, the substrate is instead printed using a conductive ink. A variety of printing processes are suitable for this technique. The conductive ink is based on small conductive particles suspended in a hardening binder. Common conductive ink materials are silver and carbon. Silver-based inks offer lower resistance compared to carbon-based inks, but are prone to surface oxitation and as such are not suitable for contacts. As such, they are often combined with a carbon ink used in contact areas. Carbon-based inks have high resistance, and can be used to create embedded resistors. The circuit shown in Figure 3.4c contains a tamper sensing mesh on a lower layer, and a keypad matrix with carbon contacts on its surface.

Figure 3.4d shows part of a mesh and a contact created using Laser Direct Structuring, an industrial technique combining selective activation of a plastic surface using a scanning laser and electroless gold plating [157]. Where in electroplating electrical current is used to deposit metal atoms

(a) Direct soldering (specimen H05).

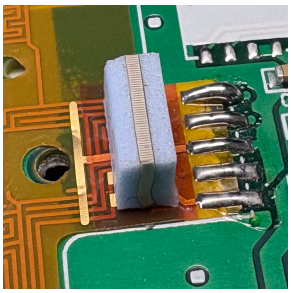

(b) Direct soldering of an FPC and an elastomeric connector (specimen H31).

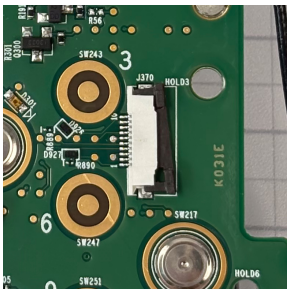

(c) Landing pads for tactile contact domes as well as FPC connector (specimen H20).

(d) Elastomeric connector landing pattern as well as stacking board-to-board connector (specimen H17).

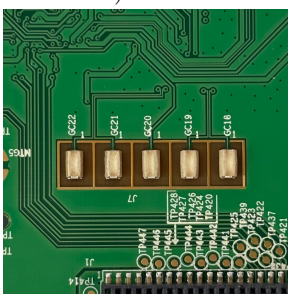

(e) Soft, conductive EM shielding gaskets used as connectors (specimen H14).

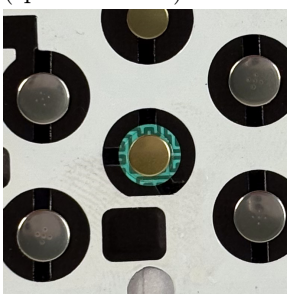

(f) Tactile dome (specimen H06).

Figure 3.5: Connecting methods used between tamper sensing mesh assemblies and their base PCBs

on a surface, in electroless plating a series of chemical reactions is used. Electroplating requires all traces to be electrically connected to form a single electrode, while electroless plating can be used on the finished circuit. Laser Direct Structuring allows patterning complex surfaces with fine structures made from metal deposited in a thin layer. In Figure 3.4d, it is visible how the trace was created using three parallel passes by the laser. The micrograph also shows the rather coarse edge structure created by LDS, which is caused by the rough surface left after pulsed laser ablation. The uneven, thin layer of metallization created by LDS results in mechanically fragile contacts that must be contacted using a soft material, usually an elastomeric connector.

Connection methods

In our survey, we found a wide variety of connecting methods used to connect tamper sensing mesh assemblies with their base PCBs with a selection

shown in Figure 3.5. Both rigid PCBs and FPCs can be soldered directly to a PCB using either a Land Grid Array (LGA) technique where pads on both PCBs are soldered facing each other, or using *castellated* edges, where pads on the base PCB are soldered sideways to holes on the top PCB that have been milled in half as shown in Figure 3.5a. FPCs can also be soldered by dragging a blob of solder across the contact as shown in Figure 3.5b, but this technique is only suitable for hand soldering. Hand soldering increases unit cost over mechanized soldering techniques such as wave soldering or reflow soldering.

FPCs are suitable for use with standard FPC connectors as shown in Figure 3.5c. These connectors mate directly to a contact area on the FPC, called *gold fingers* in industry terms. Both FPCs and rigid PCBs can be used with standard board-to-board stacking connectors such as the one visible in the center of Figure 3.5d, but their use on FPCs requires a stiffener on the FPC's back side to ensure the solder joints don't break from mechanical stress when connecting or disconnecting.

In our survey, we frequently found elastomeric connectors used to connect to both flexible and rigid tamper sensing mesh assemblies. Elastomeric connectors such as the one shown in the center of Figure 3.5b are usually used in LCD construction to contact a PCB to the LCD's Indium Tin Oxide (ITO)-coated conductive glass, but they can be used between any two parallel, conductive surfaces [15]. Elastomeric connectors consist of two insulating elastic polymer layers on the outside, with a thin strip of fine, alternating conductive and insulating elastic polymer layers sandwiched in between. In Figure 3.5b the outer insulating layers are the blue polymer, and the alternating pattern can be seen embedded in their middle. The fine alternating pattern mates to much larger pads on the two contact surfaces, ensuring that adjacent contacts are electrically insulated. In tamper sensing mesh applications, elastomeric connectors provide an intrinsic disassembly detection since they require continuous pressure to maintain electrical contact. In the top part of Figure 3.5d, a land pattern for an elastomeric connector is visible.

Elastomeric connectors are elegant and allow for multiple contacts to be made in a small area using a single elastomeric connector strip, but they are not off-the-shelf components and are always custom made to order. We found several instances where other, off-the-shelf technologies were used instead to create a pressure-sensitive connection. Figure 3.5e shows a

connection made using conductive gaskets intended for creating gapless connections between PCBs and enclosures to shield Electromagnetic Emissions (EMI). Unlike elastomeric connectors, they are not anisotropic and thus they must be cut into pieces to maintain isolation between adjacent pads. This results in a much larger contact pitch compared to other solutions.

Figure 3.5f shows another technique, here used to connect the mesh layer embedded into a key pad to a base PCB. Here, a tactile metal dome intended to be used for creating buttons in low-profile keypads is used to connect the mesh to the base PCB.

An alternative to soldering and elastomeric connectors that we did not observe during our survey but that deserves mention here is Anisotropic Conductive Film (ACF) [113]. Similar to elastomeric connectors, ACF is industrially used to contact flexible PCBs to ITO-coated glass in TFT displays. ACF comes as a double-sided tape that is bonded using pressure and sometimes high temperatures, and creates a connection between conductive surfaces on both sides of the tape. This connection has an anisotropic nature, meaning that the tape only electrically conducts from one face to the other, and not laterally. Technically, this is achieved by embedding a large number of tiny conductive spheres inside the tape that when the tape is mounted get squished between the two contact surfaces. During ACF manufacturing, the distribution of these spheres is carefully controlled to provide a reliable connection while guaranteeing adjacent spheres never touch each other.

3D construction.

While practical meshes are almost always manufactured in planar processes first, their applications usually require at least partially covering a three-dimensional volume. In our survey, we saw a number of methods being used to create three-dimensional structures from planar meshes. Figure 3.6 a-d show the major construction styles we saw among our samples. Figure 3.6a and Figure 3.6b have meshes produced as flexible printed circuits, in Figure 3.6a using a standard photolithographic copper/polyimide FPC process usually used for flexible PCBs, and in Figure 3.6b using a standard silver ink screenprinting process. The choice in Figure 3.6b not to overlap the mesh in the corner is likely caused by manufacturing considerations, since it might be difficult to ensure proper folding of a small foil tab with adhesive pre-applied.

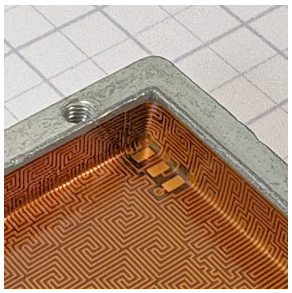

(a) Folded with overlap (specimen H03)

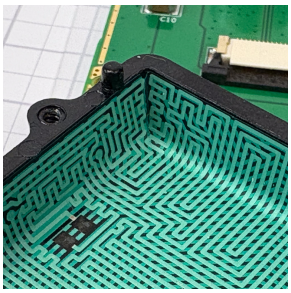

(b) Folded without overlap (specimen H14)

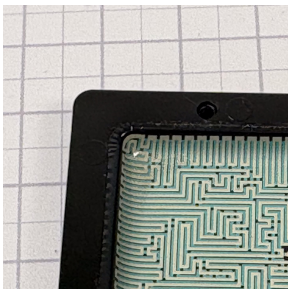

(c) Thermoformed (specimen H12)

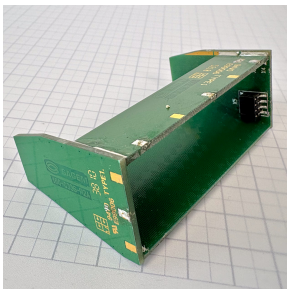

(d) House-of-Cards construction (specimen H08)

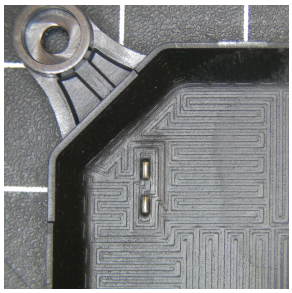

(e) Laser Direct Structuring (specimen H32)

Figure 3.6: Construction styles used to fit tamper sensing meshes into 3D envelopes. Grids in the background are 10 mm, subdivisions are 5 mm.

Figure 3.6c shows a sample of a flexible circuit manufactured in a screenprinted silver-ink process thermoformed into a three-dimensional shape [P271]. The flexible circuit mesh is first produced in a standard planar printing process. After printing and curing, the resulting foil is then heated to soften it, and forced into a three-dimensional shape using a mold. Depending on the process, one or two molds, and vacuum or pressured air can be used to shape the foil. The process requires a screenprinted flexible circuit, and would not work with copper/polyimide flexible PCBs since their copper layer is too thick to plastically deform without tearing, and because polyimide is not sufficiently thermoplastic at low temperatures.

Thermoforming is a cheap industry standard process, but applied to flexible circuits it has some limitations. First, only 2.5-dimensional structures can be created since the starting product is always a planar sheet. Second, the sheet cannot be cut or contain slots or large holes before forming since it needs to be kept under a constant tension from all sides to ensure it evenly stretches into the mold. Finally, the depth achievable in such a process is rather limited, with no sample in our survey exceeding 2 mm. Higher depths would require extensive deformation of the mesh circuit's plastic substrate,

To do
Get proper number

which could lead to tears in the mesh traces since the particle-based conductive inks used for screen-printed electronics are inelastic. Among our samples, we saw two instances of thermoformed meshes. First, all recent Ingenico terminals (H06,H13,H23,H24) integrated an ink printed mesh with thermoformed cavities into their key pad overlay. These terminals implement their key pad using tactile domes with contacts patterned on their main PCBs' surface. These domes are commonly placed on an adhesive sheet that is die cut to size so that the whole sheet can be placed on the PCB in one assembly step, instead of individually placing each dome. In these samples, a mesh was integrated into this adhesive sheet using a silver ink printing process, and two additional domes were used to provide contact between this integrated mesh and the main PCB. Cavities were formed into this mesh to enclose the upper side of the main cryptographic processor and associated components.

Figure 3.7 shows the mesh of specimen H24 both before and after removing the black opaque cover lacquer used on the bottom side of these meshes to obscure their features. The lacquer was removed by gently rubbing it with a cotton swap soaked with acetone. In Figure 3.7b, we see how the mesh's structure was adapted around the formed cavities to reduce the risk of a break during the forming process: The mesh's traces were kept parallel to the direction the foil was stretched, and the feature size of the mesh was increased by a large factor in these areas. In the corners of the formed cavity, where the foil experiences stretching in both directions, the features were scaled even larger than along the cavity's edges. This increase in structure size compromises the mesh's security level, especially given that the edges of the cavity are at a convenient direction for access by probes.

Specimen H12, shown in Figure 3.6c, displays one further design defect. The mesh shown does not extend to the edges of the plastic cover it has been molded into. When this cover is placed on top of a PCB to protect components on the PCB from tampering, this leaves a large gap between the bottom edge of the mesh and the PCB surface, through which probes can be inserted to access either the payload circuit or the mesh monitoring circuitry.

A similar design defect was mitigated in the specimens manufactured by Banksys, card payment terminal H08 and ATM encrypting pin pads H03 and H04. These specimens all have a polyimide/copper FPC mesh glued to the inside of a casted zinc lid form five sides of a cuboid. These meshes sit

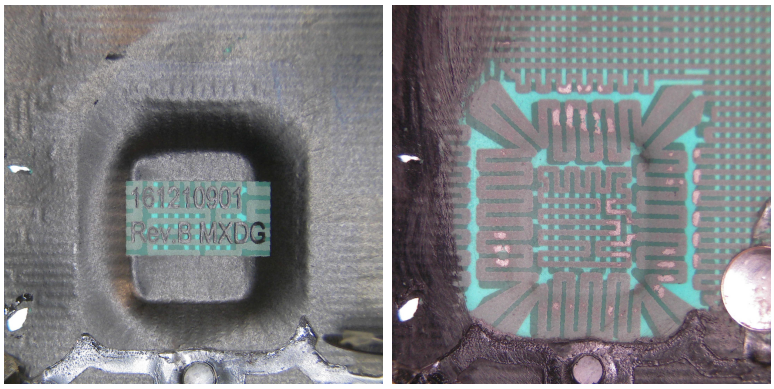

(a) Before removing opaque cover lacquer.

(b) After removing opaque cover lacquer.

Figure 3.7: Formed cavities in printed foil mesh in specimen H24.

atop their base PCBs, and a possible vulnerability would be the interface between the mesh and the PCB, where there will be an unavoidable gap of at least several hundred micrometers. In specimen H03, this was mitigated by milling a slot into the base PCB for the mesh to sit inside, thereby placing the top layer of the base PCB as well as any internal mesh layers inside the cavity of the mesh lid. In specimen H04, the payload circuit was instead placed on a daughterboard sitting inside the lid using board-to-board stacking connectors (cf. Figure 3.5d). Here, an additional rigid mesh PCB was soldered flat on top of the base PCB to cover the open side of the mesh lid, creating an overlap at the edges. In specimen H08, a card payment terminal, a simpler construction was used with a simple metal ring soldered to the base PCB mechanically shielding the edge. We are unable to ascertain why this purely mechanical shielding technique was used instead of the more secure overlapping technique seen in sample H03, which should have a similar, low manufacturing cost.

Figure 3.6e shows the result of Laser Direct Structuring (LDS), a process that avoids some of the limitations of thermoformed planar meshes. In LDS, a plastic part is covered in a conductive pattern in a combination of selective laser erosion of its surface and a series of preparation and electroless metal plating steps. LDS allows covering complex three-dimensional shapes, with the main limitation being that all patterned areas must have a direct line of sight to the outside for the scanning laser to reach it. Thus, the outside of complex parts can be covered, but internal cavities cannot. LDS is commonly used to create complex antenna shapes on the surface of internal

structural plastic parts for smartphones, but is more costly compared to screenprinting processes due to its complexity. A further disadvantage of LDS is that it is only suitable for single-layer patterns, while two layers are easily achievable in silkscreen and photolithographic PCB processes by patterning both sides of the substrate. More layers can be achived in these processes by simply stacking multiple foil layers and adding vias (through contacts), or by folding.

Figure 3.6d shows an assembly of several rigid PCBs assembled into a three-dimensional structure to protect a card slot. Solder connections between large pads are used to mechanically and electrically join the boards. While the rigid PCBs used in such as structure can be produced in a highly inexpensive, standard process, this style of construction requires manual assembly leading to increased labor cost. Furthermore, the construction leaves large gaps at edges and corners, which is not a problem for card slot protection in payment applications but which would be a flaw in a more standard HSM application.

Besides the house of cards construction style shown in Figure 3.6d where PCBs are hand-assembled into a 3D shape, rigid PCBs are also often soldered planar on top of other PCBs to serve as meshes. Figure 3.8 shows examples of such sandwich-style constructions. Figure 3.8a and Figure 3.8b show a widely used construction technique where a small mesh PCB coupon is soldered using a Land Grid Array (LGA)-technique on top of a larger base PCB containing circuitry. The goal in this technique is to project a small part of the mesh into the space above the base PCB. While this does not prevent targeted drilling as the small coupon is easy to avoid, it does prevent an attacker from sawing or laser-cutting into the side of the device parallel to the base PCB. In the implementation shown in Figure 3.8a, the coupon simply contains a small mesh embedded in an inner layer. Figure 3.8b shows a different technique, where the mesh inside the coupon is not primarily laid out in the PCB plane, but instead a large number of vias is used to create a three-dimensional zig-zag trace structure. While due to structure size limitations this via structure is much coarser than a planar mesh like that in Figure 3.8a would be, it increases the fraction of the vertical space inside the coupon that is covered by the mesh.

Figure 3.8c shows a variation of this coupon technique where two such coupons are stacked to create a small overhang, here attempting to protect the back side of a magnetic stripe reader contact in a payment terminal.

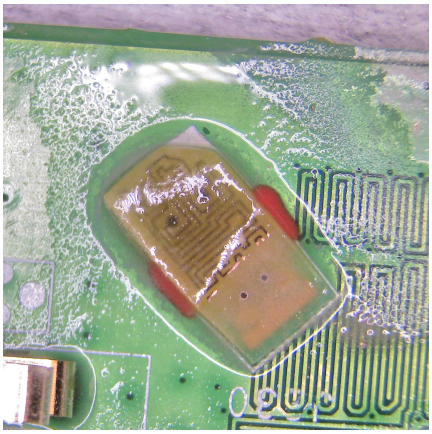

(a) Small obstacle mesh coupons (specimen H17).

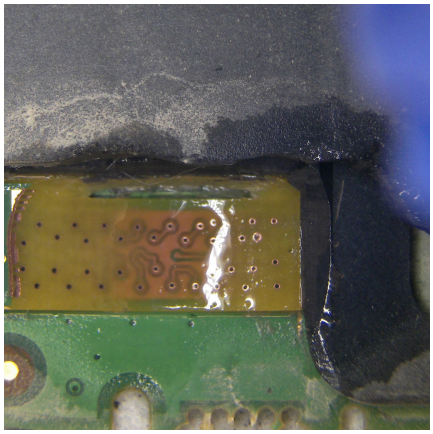

(b) Via-fence meshes (specimen H24).

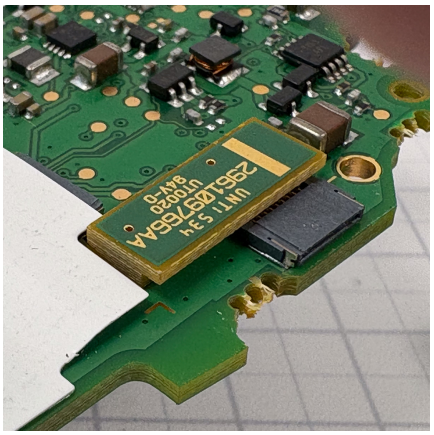

(c) Planar sandwich stack protecting the back of a connector (specimen H24).

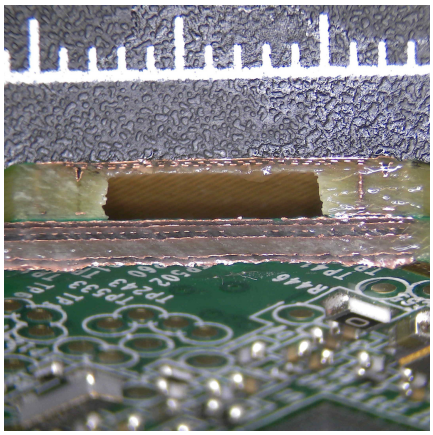

(d) PCB lid with routed cavity and embedded planar and via-fence meshes (specimen H14).

Figure 3.8: Construction styles used to cover 3D volumes using sandwich-style construction.

While a similar result could also be achieved by milling a slot into the side of a single custom-thickness PCB, the economics of PCB manufacturing are such that it may be more cost-effective to bond two standard-thickness PCBs on top of one another instead.

Figure 3.8d shows an advanced construction technique that uses a custom PCB with a large indent milled into its underside soldered on top of a base PCB to create a protected cavity on top of the base PCB. This PCB lid shows a complex internal structure. It is built up in a custom stackup with a total of six layers: A ground plane filling the top layer, then two orthogonal planar mesh layers covering the inside of the lid above the cavity. Below this standard mesh stackup are two that are used to create a via fence structure similar to that shown in Figure 3.8b in an attempt to protect the sides around the central cavity. Below these two via fence layers, at the bottom of the PCB is one more layer containing the pads connecting it to the base PCB.

CT Imaging

Hardware manufacturers implementing security meshes often attempt to keep the meshes' layouts hidden as a way of security by obscurity. In practice, this can take the form of opaque potting compounds (cf. Figure 3.9c), opaque cover layers (cf. Figure 3.4d), and burying the mesh beneath other features such as PCB ground planes (cf. Figure 3.8d, e.g. specimens H03, H17 and H32). To circumvent such attempts, an obvious attack vector is to use radiographical imaging techniques such as X-ray or CT imaging. To evaluate CT imaging as an attack method, we experimentally imaged the potted HSM module of specimen H18, an Ingenico payment terminal, using an industrial CT. Figure 3.9 shows the module we analyzed and two images exported from the resulting CT scan data. Figure 3.9a shows a horizontal cut across part of the module. In this cut, we can clearly identify a mesh layer with multiple traces, four solid metal contacts crimped to the mesh foil, and two unused contact pads and mesh traces in the lower part of the picture. An attacker would be able to use this information to target the metal contacts with a tool like a needle probe. From the CT scan we were able to measure that the mesh of the device has a pitch of 1.0 mm. Thus, even inserting a thin needle probe right through one of the mesh's traces should be possible without breaking the trace.

Figure 3.9b shows a 3D reconstruction of the mesh's conductor layout.

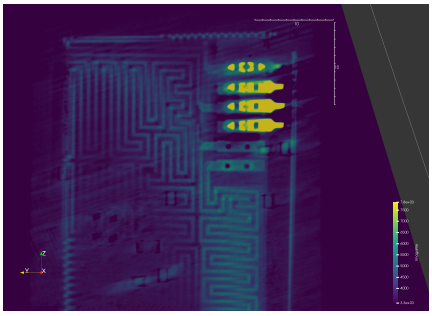

(a) CT section cut with part of a mesh layer and the crimped metal mesh contacts visible.

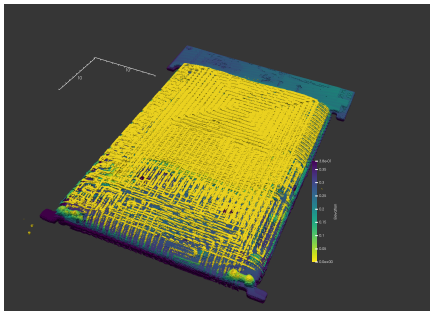

(b) CT 3D reconstruction of the mesh's trace geometry.

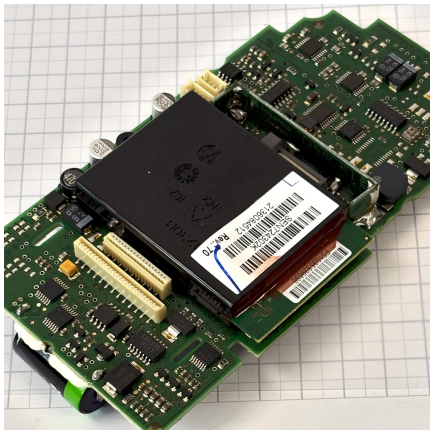

(c) Photo of the HSM module seated on the payment terminal's main PCB.

Figure 3.9: Optical photograph and CT pictures of a potted HSM module (specimen H18).

While the reconstruction is slightly noisy due to the limited scan time available, it contains ample detail to reconstruct the mesh's layout and conductor count, and even to derive conductor dimensions in order to calculate resistance and other electronic parameters. The mesh's foil is wrapped around the circuit board forming a pillow shape, which is clearly reflected in the reconstructed 3D mesh geometry. This information could be used to guide a CNC milling machine to selectively ablate the device's potting precisely down to the mesh's conductors to enable direct patching attacks on the mesh.

Results summary

Below is a table representing which features discussed in the sections above we found in which of our samples. Overall, we commonly found a combination of a rigid PCB mesh in the specimen's main PCB and and flexible meshes formed into a lid structure above its main PCB. The mesh inside the rigid PCB would protect the payload components soldered to the top surface of the PCB such as pin pad buttons or crytographic coprocessors from probing from underneath, while the flexible mesh lid would protect them from attacks from above or from the side. We only found two specimens that wrapped an entire payload PCB inside of a mesh, the Utimaco datacenter HSM appliance (H30) and an older Ingenico payment terminal (H18). Only the datacenter HSM followed this approach through, its manufacturer going to some length to carefully fold the mesh around corners and the entry point of its Flat Flex Cable (FFC) connections to the outside world to avoid possible weak points there. The payment terminal module had weak points at the corners of the wrapped mesh, and its wrapping pattern only covered five of the six sides of a cuboid, with the remaining side left open to allow for the payload PCB to pass out of the mesh for its external connections.

We found an approximately even split between copper/polyimide FPCs and silver ink printing processes being used for flexible meshes. Printed carbon ink processes were less popular, presumably because they offer no significant cost savings but the resulting mesh has a much higher electrical resistance, limiting possible mesh length.

We found potting was only infrequently used across our sample, presumably because of the limited protection it provides. We found conductive ink printed meshes commonly used opaque base foils and opaque lacquer

cover layers to obscure their features, but when dissecting these specimens we noticed that usually these opaque lacquers are easily removed without damaging the underlying printed mesh traces using a cotton swab soaked in acetone. Additionally, in almost all instances the trace structure was easily recognizable from the mesh traces' thickness showing through to the surface of the opaque cover lacquer. In practice it served as electrical insulation, but did not convey meaningful protection against reverse engineering.

Feature	Figures	Specimen 1	2	3	4	5	6	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	27	28	30	31	32
Mesh Contacts.																														
Elastomeric	3.5b, 3.5d	●	●	●	●			●		●		●		●	●	●	●			●	●	●				●	●		●	●
Soldered	3.5a	●			●	●	●	●				●				●	●	●				●	●	●					●	
Stacking	3.5d							●									●													
Tactile Dome	3.5f, 3.5c						●						●								●		●	●						
FPC Connector	3.5c					●			●	●				●	●	●	●	●			●							●		
Mesh EMI Gasket	3.5e													●																
Mesh Material																														
Rigid PCB	3.4a	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●			●	●	●	●		●	●		●	●
Copper FPC	3.4b			●	●		●	●				●	●			●	●			●	●		●	●		●	●		●	
Printed silver ink	3.4c					●	●			●		●	●	●			●	●			●		●	●						
Printed carbon ink	3.4e	●							●																			●		
Gold (Laser Direct Structuring)	3.4d																													●
3D Construction																														
Folded mesh	3.6a, 3.6b	●		●	●	●	●	●	●	●		●	●			●	●	●		●	●		●	●		●		●	●	
House of cards	3.6d	●						●		●								●												
Laser Direct Structuring	3.6e																													●
Thermoformed	3.6c, 3.7						●					●											●	●						
Planar obstacle	3.8a, 3.8b	●			●	●							●	●		●	●						●							
Complex planar	3.8c, 3.8d				●									●																
Obscurity Features																														
Metal enclosure	3.6a		●	●	●			●								●					●					●		●		
Potting	3.9c					●												●										●		
Opaque foil	3.5f					●	●		●	●			●				●			●	●		●					●		
Opaque lacquer	3.7					●	●						●				●			●			●					●		●
Other Features																														
Integrated tactile domes	3.5f					●	●						●				●				●		●	●		●			●	
Integrated contact pads	3.5c													●		●					●					●				

Table 3.2: Feature matrix of all specimens analyzed. Dots indicate presence of a feature. The figures column lists which figures above contain examples of a particular feature.

4 Discussion

In our survey, we have seen the technological state of the art to which tamper-sensing meshes have evolved since the earliest designs evidenced in patents from 150 years ago. While mesh manufacturing technology has experienced some advancements from historical wire-wound meshes to modern meshes always being constructed in printed circuit processes, mesh monitoring approaches have received surprisingly little attention through the centuries and even in recent, state-of-the-art systems, a simple comparator monitoring a mesh arranged in a bridge configuration is still considered sufficient in high-security applications [195].

4.1 Mesh construction techniques

We found that in almost all cases, practical tamper sensing meshes are constructed using standard manufacturing processes. In some card payment terminals, we found meshes that used slightly customized standard processes and e.g. integrated a mesh layer produced in a carbon printing process into a membrane keypad, but customizations were minimal. We only found one mesh manufactured in a bespoke process in the datacenter HSM appliance we examined, and that bespoke process turns out to be a turnkey solution used by at least two HSM vendors. Underscoring stagnating development in the field, this particular mesh manufacturing process seems to have seen only minimal changes since the first patents covering it were published in the late 1990ies [P160, P161, 195].

4.2 Mesh monitoring circuits

We observed that in general, academic research leads before patent literature, which is ahead of actual implementations in the field. Practical monitoring circuitry seems basic. Particularly the datacenter HSM appliance we examined (specimen H30) showed a contrast between a mesh manufactured in a bespoke process combined with an unsophisticated, discrete monitoring circuit based around a number of voltage comparators [195]. We will go into more detail on improved monitoring methods as well as the academic state of the art in this field in Chapter 5.

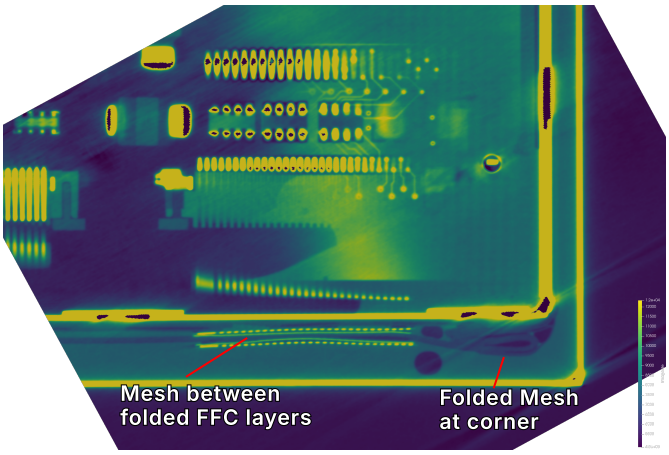

Figure 3.10: Computed Tomography (CT) scan of a corner of the PCIe HSM module from an Utimaco rackmount HSM appliance. Visible are several capacitors, the edge of a large IC, and a large Flat Flexible Cable (FFC) connector. Two layers of metal enclosures with resin potting in between are visible, and the security mesh can be seen folded between layers of the folded FFC cable connecting to the outside.

4.3 Computed Tomography Imaging

CT imaging presents a serious threat to any HSM design that relies on its mesh layout remaining secret. For instance, the Gore tamper sensing mesh product used in IBM and Utimaco HSMs includes a feature where after production, small vias are lasered into a specially prepared area on the mesh foil to randomize the connection pattern of the mesh on a unit-by-unit basis. CT imaging could be used to discern this type of customization. Furthermore, CT imaging can be used to provide sub-millimeter accurate positioning for an attack, even if the specimen to be attacked has large production tolerances. We found that CT imaging can be made more difficult using three complementary techniques.

Low-contrast trace materials. CT imaging can be made more difficult by manufacturing the mesh with very thin conductive traces, and using a trace material that has low atomic number, corresponding to low X-ray absorption. For instance, the Gore mesh specimen used a carbon-based ink that judging by structure size was screen-printed, which leads to an economical yet relatively secure solution [14, 236].

Use of X-ray attenuating materials. We found that placing any highly X-ray attenuating material in the HSM makes CT imaging more difficult.

Figure 3.10 shows a CT image taken from an Utimaco HSM. The device has two thick metal layers with a potting resin and the tamper sensing mesh in between, so high-energy X-rays were necessary to penetrate both metal layers and image the device. As a result, the contrast on X-ray-transparent features like polymers is low. In comparison, the Ingenico specimen was easy to image since it consisted of a PCB wrapped with a mesh foil and encased in resin inside of an injection-molded plastic enclosure. Thus, we were able to image it at a low X-ray energy and we were able to easily reconstruct detail on both the mesh's layout and the PCB's circuitry. To apply X-ray dense materials for defense in a practical design, a sheet made from elementary tin or a tin alloy would be a suitable choice for such an X-ray absorbing feature since tin is cheap, non-hazardous and absorbs X-rays almost as well as lead. Alternatively to a sheet-metal enclosure, an X-ray absorbing material could also be incorporated into a potting compound as a powder.

Size. Finally, we found that a larger module size makes CT imaging more difficult simply due to the thickness of material that the X-rays need to penetrate. Ideally, a HSM should aim for a cuboid form factor, as the common flat construction style is easily penetrated by X-rays along at least one axis.

Radiation sensors. Besides engineering techniques making CT imaging harder, in battery-powered devices with active tamper sensing, CT imaging can be actively detected to trigger a tamper alarm. During CT imaging, a large amount of high-energy X-ray images are taken. X-ray radiation can be reliably detected using off-the-shelf sensors that usually consist of a large-area photodiode coupled to a scintillator crystal converting X-ray photons to visible light.

5 Conclusion

In this survey, we have analyzed a wide variety in tamper sensing mesh construction techniques. Meshes are commonly implemented as part of both rigid (PCB) and flexible (FPC) circuit boards, either standalone, or as part of a board also carrying other components. Silver or carbon trace patterning techniques that are normally used for membrane keyboards are also used in some meshes, but are limited in their structure size. The meshes we found

in the wild almost never push the boundaries of achievable structure size for a given process.

The strongest systems we found combined a mesh with potting such that separating mesh and potting destroyed the mesh's traces. Silver or carbon ink printed circuits like they are normally used for keyboard matrices performed particularly well in this regard since such inks adheres better to some potting compounds than to its plastic carrier substrate. We found copper FPCs are commonly used for meshes. Interestingly, they seem to be a poor choice since they are very robust and can even be forcibly separated from some potting compounds without destroying their traces.

The weakest systems we found completely omitted a tamper sensing mesh. Ironically, all of these systems were devices marketed as hardware security modules. Given the inexpensive nature of tamper sensing meshes and the high price point of such devices, we suspect market segmentation as a driving force behind their manufacturers' decision to omit tamper sensing meshes despite their low cost. The primary security standard that is most often cited for the certification of HSMs is the US government's FIPS-140, now in its third version [2]. A peculiarity of this standard is that it only requires active tamper sensing meshes in the highest of the four security levels it defines. Overall, we can conclude that the term "HSM" does not imply state-of-the-art physical tamper sensing.

From an academic point of view, the core finding of our survey is that for academic research on mesh manufacturing, monitoring or attacks on meshese, realistic tamper sensing mesh samples can easily be created. A number of commercial manufacturing processes would yield acceptable standins for real devices found in the wild. With the exception of a single device that used a particularly fine structure size in the 100 µm range approaching the limit of inexpensive PCB manufacturing processes, none of the devices we examined utilized particularly non-obvious construction techniques.

From an engineering point of view, we observe that across application domains, tamper sensing meshes often use basic construction techniques for both the mesh itself and for its monitoring circuit. Implementing such a system that matches the security of devices seen in the wild should be achievable to most engineers.

We find that the IHSM approach is a natural extension of the state of the art that we saw reflected in tamper sensing mesh implementations in the field, and that the construction techniques that have been applied to

improve their security can be carried over to IHSM implementations. The three-dimensional layout of a mesh becomes easier in an IHSM implementation since features like corners between mesh panels or gaps between mesh layers in most layouts are protected by the mesh's motion. An unintended advantage that results in IHSM implementations over conventional meshes is that they would provide a level of intrinsic resistance to X-ray and CT imaging. In contrast to optical cameras in the visible spectrum, X-ray image sensors need integration times in the hundreds of milliseconds or longer, which makes them unsuitable to image a quickly moving target.

Web sources

[W19] Banque centrale du Luxembourg. *Ink-Stained Banknotes*. URL: https://www.bcl.lu/en/Banknotes-and-Coins/remboursement/billets-macules1/index.html (visited on 2025-11-21) (cit. on p. 43).

[W65] European Central Bank. *Damaged and Ink-Stained Banknotes*. 2023-07-10. URL: https://www.ecb.europa.eu/euro/banknotes/damaged/html/index.en.html (visited on 2025-11-21) (cit. on p. 43).

[W125] *ISO/IEC 19790:2025*. ISO. URL: https://www.iso.org/standard/82423.html (visited on 2025-05-15) (cit. on pp. 5, 24, 38).

[W144] Kruse Sicherheitssysteme. *Datenblatt KRUSE FW-Schlüsseldepot Basic*. 2018-12. URL: https://kruse-shop.de/media/pdf/e3/c0/6c/MA-KRUSE-FW-Schluesseldepot-FSD-D-E_Rev1-3-20-12-18.pdf (visited on 2025-10-30) (cit. on p. 36).

[W173] mikeselectricstuff. *Neopost Postal Franking Machines*. 2023-10-03. URL: https://www.youtube.com/watch?v=eO7AoHI2Tpk (visited on 2025-02-17) (cit. on p. 43).

[W232] Securitas Technology GmbH. *SD-04203RB25-D5*. Setec Sicherheitstechnik. 2019. URL: https://setec-security.de/wp-content/uploads/2019/11/SD-04203RB25-D5.pdf (visited on 2025-10-30) (cit. on p. 36).

[W249] Thales Group. *Thales Luna Network HSM 7 Functionality Module Software Development Kit Guide*. 2025-11-26. URL: https://thalesdocs.com/gphsm/luna/7/docs/network/Content/PDF_Network/FM%20SDK%20Programming%20Guide.pdf (visited on 2025-12-01) (cit. on p. 42).

Patent References

[P33] William L. Brodsky et al. "Circuit Layouts of Tamper-Respondent Sensors". U.S. pat. 10136519B2. International Business Machines Corp. 2018-11-20 (cit. on p. 33).

[P41] Mario Cesana and Roberto Zavatti. "Tamper Resistant Card Enclosure with Improved Intrusion Detection Circuit". U.S. pat. 20010056542A1. International Business Machines Corp. 2001-12-27 (cit. on pp. 33, 38).

[P42] Mario L. Cesana, Donald S. Farquhar, and Martino Taddei. "Security Cloth Design and Assembly". U.S. pat. 6982642B1. International Business Machines Corp. 2006-01-03 (cit. on p. 33).

[P47] Douglas A. Clark. "Tamper Detection System for Securing Data". U.S. pat. 6895509B1. Pitney Bowes Inc. 2005-05-17 (cit. on p. 33).

[P48] Cornel P. Cobianu, Ion Georgescu, and Viorel-Georgel Dumitru. "Large Area Distributed Sensor". U.S. pat. 20080001741A1. Honeywell International Inc. 2008-01-03 (cit. on p. 33).

[P49] Timothy E. Cook and Gerald Thomas Wardrop Jr. "Tamper Detection Circuit Assemblies and Related Manufacturing Processes". U.S. pat. 10579833B1. Thales eSecurity Inc. 2020-03-03 (cit. on p. 33).

[P54] Claude Société Civile S. P. I. D. Dalphin. "Enceinte Protégée Avec Interrupteur Électrique et Son Application". European pat. 0231549A1. Telecommunications Radioelectriques et Telephoniques SA TRT, Philips Gloeilampenfabrieken NV, Koninklijke Philips Electronics NV. 1987-08-12 (cit. on p. 38).

[P59] Hartmut Droege et al. "Sicherheitsmodul Mit Einteiliger Sicherheitsfolie". German pat. 19600769A1. International Business Machines Corp. 1997-07-17 (cit. on p. 33).

[P61] Arcadi Elbert and Alvin Diep. “Secure Circuit Assembly”. U.S. pat. 20060259788A1. Individual. 2006-11-16 (cit. on p. 33).

[P62] “Elektrische Sicherheitseinrichtung Zum Schutze von Geldschraenken u. Dgl”. German pat. 559905C. Individual. 1932-09-26 (cit. on p. 38).

[P101] Conrad S. Ham and Elwood R. Horwinski. “Printed-Circuit Type Security Apparatus for Protecting Areas”. U.S. pat. 3594770A. Lewis Engineering Co. 1971-07-20 (cit. on pp. 33, 38).

[P103] Kjell Heitmann, Douglas Clark, and Paul Perreault. “Tamper Barrier for Electronic Device”. U.S. pat. 20050161253A1. Pitney Bowes Inc. 2005-07-28 (cit. on p. 33).

[P104] Kjell A. Heitmann, Douglas A. Clark, and Paul G. Perreault. “Method of Making Tamper Detection Circuit for an Electronic Device”. U.S. pat. 7475474B2. Pitney Bowes Inc. 2009-01-13 (cit. on p. 33).

[P105] Maxim Hennig et al. “Apparatus and Method Comprising a Carrier with Circuit Structures”. U.S. pat. 14867889. Fraunhofer Gesellschaft zur Foerderung der Angewandten Forschung eV. 2020-03-17 (cit. on p. 37).

[P120] “Improvement in Electro-Magnetic Envelopes for Safes, Vaults”. U.S. pat. 110362A. 1870-12-20 (cit. on p. 33).

[P121] “Improvement in Protecting Safes and Vaults from Burglars”. U.S. pat. 106324A. 1870-08-16 (cit. on pp. 33, 37).

[P129] Richard J. Joyce and Allan R. Kramer. “Method to Detect Penetration of a Surface and Apparatus Implementing Same”. U.S. pat. 5568124A. Hughes Aircraft Co. 1996-10-22 (cit. on p. 33).

[P136] Theodoor A. Kleijne. “Security Device for the Secure Storage of Sensitive Data”. U.S. pat. 4593384A. NCR Corp. 1986-06-03 (cit. on p. 33).

[P160] Hugh MacPherson. “Tamper Respondent Enclosure”. U.S. pat. 5858500A. WL Gore and Associates Inc. 1999-01-12 (cit. on pp. 37, 63).

[P161] Hugh Macpherson. “Improvements in Security Enclosures”. European pat. 0540139A2. WL Gore and Associates UK Ltd. 1993-05-05 (cit. on pp. 37, 63).

[P196] Johannes Obermaier, Vincent Immler, and Robert Hesselbarth. "PUF-film and Method for Producing the Same". U.S. pat. 11586780B2. Fraunhofer Gesellschaft zur Foerderung der Angewandten Forschung eV. 2023-02-21 (cit. on pp. 3, 37).

[P204] Paul Perreault, Douglas Clark, and Kjell Heitmann. "System and Method for Installing a Tamper Barrier Wrap in a PCB Assembly, Including a PCB Assembly Having Improved Heat Sinking". U.S. pat. 20050160702A1. Pitney Bowes Inc. 2005-07-28 (cit. on p. 33).

[P207] Cuong V. Pham et al. "Anti-Tamper Mesh". U.S. pat. 7947911B1. Teledyne Technologies Inc. 2011-05-24 (cit. on p. 33).

[P215] Mani Razaghi. "Circuit Board to Hold Connector Pieces for Tamper Detection Circuit". U.S. pat. 10251260B1. Square Inc. 2019-04-02 (cit. on pp. 33, 38).

[P244] Henry M. Sutton, Walter L. Steele, and Michael Coerver. "Electrically-Protected Structure". U.S. pat. 708093A. Individual. 1902-09-02 (cit. on pp. 33, 38).

[P271] Karl Weidner and Anton Wimmer. "Hardwareschutz in Form von zu Halbschalen tiefgezogenen Leiterplatten". Pat. WO2007003227A1 (WO). Siemens Aktiengesellschaft. 2007-01-11 (cit. on p. 53).

References

[2] (US) National Institute of Standards and Technology. *Security Requirements for Cryptographic Modules*. Federal Information Processing Standard (FIPS) 140-3. U.S. Department of Commerce, 2019-03-22. DOI: 10.6028/NIST.FIPS.140-3 (cit. on pp. 2, 4, 24, 38, 66).

[14] Ross Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 1st ed. Wiley, 2020-12-22. DOI: 10.1002/9781119644682 (cit. on pp. 2, 22, 24, 36–38, 43, 64, 75, 77, 78, 92, 118).

[15] Davide Andrea. *The Electronic Connector Book*. 1st ed. 2022. ISBN: 978-1-300-09248-3 (cit. on p. 51).

[26] Barry M. Blechman, ed. *Technology and the Limitation of International Conflict*. FPI Papers in International Affairs. Washington, DC: Foreign Policy Inst. [u.a.], 1989. 185 pp. ISBN: 978-0-941700-42-9 978-0-941700-43-6 (cit. on p. 31).

[28] David G. Boak. *A History of U.S. Communications Security (The David G. Boak Lectures), Volume I.* (US) National Security Agency, 1973 (cit. on p. 33).

[29] David G. Boak. *A History of U.S. Communications Security (The David G. Boak Lectures), Volume II.* (US) National Security Agency, 1981 (cit. on p. 33).

[39] Ashton Carter et al., eds. *Managing Nuclear Operations*. Washington, D.C: Brookings Institution, 1987. 751 pp. ISBN: 978-0-8157-1313-5 978-0-8157-1314-2 (cit. on p. 34).

[109] Paul Horowitz and Winfield Hill. *The Art of Electronics*. Third edition, 21st printing with corrections. Cambridge, New York: Cambridge University Press, 2024. 1230 pp. ISBN: 978-0-521-80926-9 (cit. on p. 38).

[113] Andrew "bunnie" Huang. *The Hardware Hacker: Adventures in Making and Breaking Hardware*. San Francisco: No Starch Press, 2019. 1 p. ISBN: 978-1-59327-758-1 978-1-59327-813-7 (cit. on p. 52).

[116] 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)*. 2018 IEEE International Symposium on Hardware Oriented Security and Trust (HOST). 2018-04, pp. 49–56. DOI: 10.1109/HST.2018.8383890 (cit. on pp. 3, 37, 115, 116, 118, 124).

[117] Vincent Immler et al. "Secure Physical Enclosures from Covers with Tamper-Resistance". In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2018-11-09), pp. 51–96. DOI: 10.46586/tches.v2019.i1.51-96 (cit. on pp. 37, 115, 141, 198).

[122] International Atomic Energy Agency. *Safeguards, Techniques and Equipment*. Vol. 1. International Nuclear Verification Series. 2011. ISBN: 978-92-0-118910-3 (cit. on pp. 35, 36, 78).

[157] LPKF Laser & Electronics AG. *LPKF LDS: Laser Direct Structuring for 3D Molded Interconnect Devices*. 2014 (cit. on p. 49).

[195] Johannes Obermaier and Vincent Immler. "The Past, Present, and Future of Physical Security Enclosures: From Battery-Backed Monitoring to PUF-based Inherent Security and Beyond". In: *Journal of Hardware and Systems Security* 2 (2018), pp. 289–296. DOI: 10.1007/s41635-018-0045-2 (cit. on pp. 2, 37, 63, 78, 218).

[198] Oberthur Cash Protection. *Introduction to Cash Protection: Intelligent Banknote Neutralization Systems*. 2019 (cit. on p. 43).

[202] PCI Security Standards Council. *Payment Card Industry PIN Transaction Security Device Testing and Approval Program Guide*. Version 2.2. 2025-06 (cit. on p. 39).

[234] G.J. Simmons. "How to Insure That Data Acquired to Verify Treaty Compliance Are Trustworthy". In: *Proceedings of the IEEE* 76.5 (1988-05), pp. 621–627. DOI: 10.1109/5.4446 (cit. on p. 35).

[236] Sean W Smith and Steve Weingart. "Building a High-Performance, Programmable Secure Coprocessor". In: *Computer Networks* 31.8 (1999-04), pp. 831–860. DOI: 10.1016/S1389-1286(98)00019-X (cit. on pp. 37, 64, 78, 95).

[252] Keith Tolk et al. "Safeguards Sensors and Systems: Past, Present, and Future". In: *Journal of Nuclear Materials Management* 35.4 (2007-07-01), pp. 101–110. ISSN: 0893-6188 (cit. on p. 36).

[264] 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). 2019-10, pp. 212–215. DOI: 10.1109/SIITME47687.2019.8990877 (cit. on pp. 37, 115, 118, 119).

Chapter 4

Inertial Hardware Security Modules

Otiluke's Resilient Sphere

Level 4 Abjuration (Wizard)

Casting Time: Action
Range: 30 feet
Components: V, S, M (a glass sphere)
Duration: Concentration, up to 1 minute

A shimmering sphere encloses a Large or smaller creature or object within range. An unwilling creature must succeed on a Dexterity saving throw or be enclosed for the duration.

Nothing—not physical objects, energy, or other spell effects—can pass through the barrier, in or out, though a creature in the sphere can breathe there. The sphere is immune to all damage, and a creature or object inside can't be damaged by attacks or effects originating from outside, nor can a creature inside the sphere damage anything outside it.

The sphere is weightless and just large enough to contain the creature or object inside. An enclosed creature can take an action to push against the sphere's walls and thus roll the sphere at up to half the creature's Speed. Similarly, the globe can be picked up and moved by other creatures.

A *Disintegrate* spell targeting the globe destroys it without harming anything inside.

Contents

1 Introduction

This part is adapted from a paper written by me and presented by me at CHES 2022 [94].

While information security technology has matured a great deal in the last half-century, physical security did not keep up with the pace of the remainder of this industry. Given the right skills, physical access to a computer still often allows full compromise. The physical security of modern server hardware hinges on what lock you put on the room it is in.

Currently, servers and other computers are rarely physically secured as a whole. Servers sometimes have a simple lid switch and are put in locked "cages" inside guarded facilities. This usually provides a good compromise between physical security and ease of maintenance. To handle highly sensitive data in applications such as banking or public key infrastructure, general-purpose and low-security servers are augmented with dedicated, physically secure cryptographic co-processors such as trusted platform modules (TPMs) or hardware security modules (HSMs). Using a limited amount of trust in components such as the CPU, the larger system's security can then be reduced to that of its physically secured TPM [W187, 73, 128]. Like smartcards, TPMs rely on a modern IC being hard to tamper with. Shrinking things to the nanoscopic level to secure them against tampering is a good engineering solution for some years to come. However, in essence, this is a type of security by obscurity: Obscurity here referring to the rarity of the equipment necessary to attack modern ICs [8, 14].

In contrast to TPMs and Smartcards, HSMs rely on an active security barrier usually consisting of a fragile foil with conductive traces. These traces are much larger scale than a smart card IC's microscopic structures and instead are designed to be very hard to remove intact. While we are certain that there still are many insights to be gained in both technologies, we wish to introduce a novel approach to sidestep the manufacturing issues of both and provide radically better security against physical attacks. Our core observation is that any cheap but coarse HSM technology can be made much more difficult to attack by moving it very quickly.

For example, consider an HSM as it is used in online credit card payment processing. Its physical security level is set by the structure size of its security mesh. An attack on its mesh might involve fine drill bits, needles, wires, glue, solder, and lasers [58]. Now consider the same HSM mounted on a large flywheel. In addition to its usual defenses, this modified HSM is now equipped with an accelerometer that it uses to verify that it is spinning at high speed. How would an attacker approach this HSM? They would have

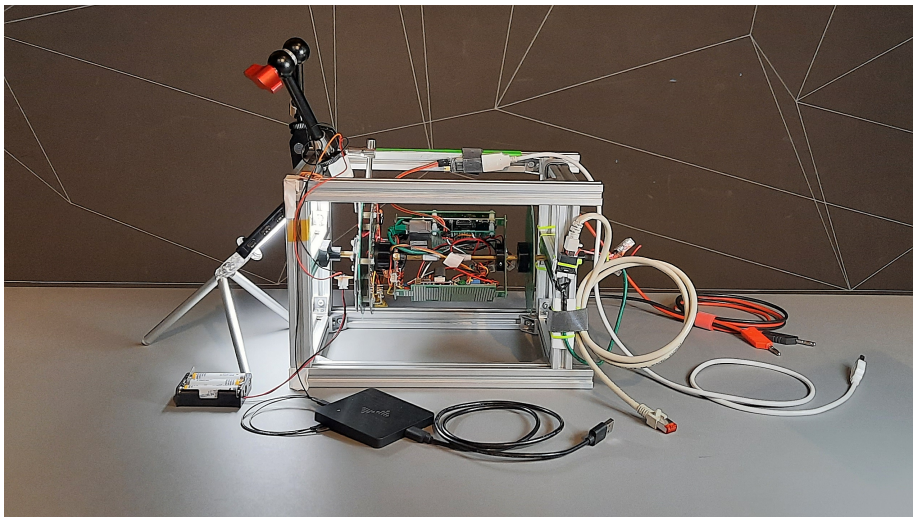

Figure 4.1: The prototype as we used it to test power transfer and bidirectional communication between stator and rotor. This picture shows the proof-of-concept prototype's configuration that we used for accelerometer characterization (Section 6) without the vertical security mesh struts that connect the circular top and bottom outer meshes.

to either slow down the rotation—which triggers the accelerometer's monitoring circuit—or they would have to attack the HSM in motion. The HSM literally becomes a moving target. At slow speeds, rotating the entire attack workbench might be possible—but rotating frames of reference quickly become inhospitable to human life (see Section 4.3). Since non-contact electromagnetic or optical attacks are more limited in the first place and can be shielded, we have effectively forced the attacker to use an "attack robot".

This chapter contains the following contributions:

1. We present the *Inertial HSM* concept. Inertial HSMs enable cost-effective, small-scale production of highly secure HSMs.

2. We discuss possible tamper sensors for inertial HSMs.

3. We explore the design space of our inertial HSM concept.

4. We present our work on a prototype inertial HSM (Figure 4.1).

5. We present an analysis of the viability of using commodity MEMS accelerometers as braking sensors.

In Section 2, we will give an overview of the state of the art in HSM physical security. On this basis, in Section 3 we will elaborate the principles of our Inertial HSM approach. We will analyze its weaknesses in Section 4.

Based on these results we have built a proof-of-concept hardware prototype. In Section 5 we will elaborate on the design of this prototype. In Section 6 we present our characterization of an automotive MEMS accelerometer IC as a rotation sensor in this proof-of-concept prototype. We conclude this chapter with a general evaluation of our design in Section 6.

2 Related work

In this section, we will briefly explore the history of HSMs and the state of academic research on active tamper detection.

HSMs are an old technology that traces back decades in its electronic realization, initially being conceived by the US NSA during the second world war [30]. Today's common approach of monitoring meandering electrical traces on a fragile foil that is wrapped around the HSM essentially transforms the security problem into the challenge to manufacture very fine electrical traces on a flexible foil [123, 118, 14]. There has been some research on monitoring the HSM's interior using e.g. electromagnetic radiation [251, 143] or ultrasound [266] but none of this research has found widespread adoption yet.

HSMs can be compared to physical seals [14]. Both are tamper-evident devices. The difference is that an HSM continuously monitors itself whereas a physical seal only serves to record tampering and requires someone to examine it. This examination can be done by eye in the field, but it can also be carried out in a laboratory using complex equipment. An HSM in principle has to have this examination equipment built-in.

Physical seals are used in a wide variety of applications. The most interesting ones from a research point of view that are recorded in public literature are those used for the monitoring of nuclear material under the International Atomic Energy Authority (IAEA). Most of these seals use the same approach that is used in Physical Unclonable Functions (PUFs), though their development predates that of PUFs by several decades. The seal is created in a way that intentionally causes large, random device-to-device variations. These variations are precisely recorded at deployment. At the end of the seal's lifetime, the seal is returned to a lab and closely examined to check for any deviations from the seal's prior recorded state. The type of variation used in these seals includes random scratches in metal parts and random blobs of solder (IAEA metal cap seal), randomly cut op-

tical fibers (COBRA seal), the uncontrollably random distribution of glitter particles in a polymer matrix (COBRA seal prototypes) as well as the precise three-dimensional surface structure of metal parts at microscopic scales (LMCV) [122].

The IAEA's equipment portfolio does include electronic seals such as the EOSS. These devices are intended for remote reading, similar to an HSM. They are constructed from two components: A cable that is surveilled for tampering, and a monitoring device. The monitoring device itself is in effect an HSM and uses a security mesh foil like it is used in commercial HSMs.

The self-destruct built into an HSM serves as a strong tamper deterrent. For illustration, compare an HSM to a computer inside a locked safe when opposing a well-funded attacker with plenty of time. In [30], Boak asserts that absent an HSM's capability to self-destruct, the best safes can only withstand brute force attacks by an expert for several minutes. While the state of electronics has advanced rapidly since Boak's 1973 lecture, the hardness of steel has not increased correspondingly. Thus, we can conclude that even today, against a "smart, well-equipped opponent with plenty of time" as noted by Boak, this self-destruction functionality is essential.

In [14], Anderson gives a comprehensive overview of physical security. An example HSM that he cites is the IBM 4758, the details of which are laid out in-depth in [236]. This HSM is an example of an industry-standard construction. Although its turn of the century design is now a bit dated, the construction techniques of the physical security mechanisms have not evolved much in the last two decades. Besides some auxiliary temperature and radiation sensors to guard against attacks on the built-in SRAM memory, the module's main security barrier uses the common construction of a flexible mesh foil wrapped around the module's core. In [236], the authors state that the module monitors this mesh for short circuits, open circuits, and conductivity. Other commercial offerings use similar approaches to tamper detection [195, 58, 14, 123].

Shifting our focus from industry use to the academic state of the art, in [118], Immler et al. describe an HSM based on precise capacitance measurements of a security mesh, creating a PUF from the mesh. In contrast to traditional meshes, they use a large number of individual traces (more than 30 in their example). Their concept promises a very high degree of protection but is limited in the board area covered and component height, as well as the high cost of the advanced analog circuitry required for mon-

itoring. A core component of their design is that they propose its use as a PUF to allow for protection even when powered off, similar to a smart card—but the design is not limited to this use.

In [251], Tobisch et al. describe a construction technique for a hardware security module that is based on a WiFi transceiver inside a conductive enclosure. In their design, a reference signal is sent into the RF cavity formed by the conductive enclosure. One or more receivers listen for the signal's reflections and use them to characterize the phase and frequency response of the RF cavity. The assumption underlying their system is that the RF behavior of the cavity is inscrutable from the outside and that any small disturbances within the volume of the cavity will cause a significant change in its RF response. A core component of the work of Tobisch et al. [251] is that they use commodity WiFi hardware, so the resulting system is likely both much cheaper and capable of protecting a much larger security envelope than designs using finely patterned foil security meshes such as [118], at the cost of worse and less predictable security guarantees. Where [251] use electromagnetic radiation, Vrijaldenhoven in [266] uses ultrasound waves traveling on a surface acoustic wave (SAW) device to a similar end.

While Tobisch et al. [251] approach the sensing frontend cost as their primary optimization target, the prior work of Kreft and Adi [143] considers sensing quality. Their target is an HSM that envelopes a volume barely larger than a single chip. They theorize how an array of distributed RF transceivers can measure the physical properties of a potting compound that has been loaded with RF-reflective grains. In their concept, the RF response characterized by these transceivers is shaped by the precise three-dimensional distribution of RF-reflective grains within the potting compound.

To the best of our knowledge, we are the first to propose a mechanically moving security barrier as part of a hardware security module. Most academic research concentrates on the issue of creating new, more sensitive security barriers for HSMs [118] while commercial vendors concentrate on means to certify and cheaply manufacture these security barriers [58]. Our concept instead focuses on the issue of taking any existing, cheap low-performance security barrier and transforming it into a marginally more expensive but high-performance one. The closest to a mechanical HSM that we were able to find during our research is a 1988 patent [P214] that describes a mechanism to detect tampering along a communication cable

by enclosing the cable inside a conduit filled with pressurized gas.

3 Inertial HSM construction and operation

Fast mechanical motion has been proposed as a means of making things harder to see with the human eye [W100] and is routinely used in military applications to make things harder to hit [W247] but we seem to be the first to use it in tamper detection.

The core questions in the design of an inertial HSM are the following:

1. What **type of motion** to use, such as rotation, pendulum motion, or linear motion.
2. How to construct the **tamper detection sensor**.
3. How to **detect braking** of the IHSM's movement.
4. The **mechanical layout** of the system.

We will approach these questions one by one in the following subsections and conclude this section with an exploration of the practical implications that these aspects of IHSM construction have on IHSM operation, but first, we will motivate our concept with two use cases and outline our attacker model.

3.1 Use Cases and Attacker Model

The target application of an IHSM is high-risk data processing. This risk can be implied by either high-value data, or by difficult physical security constraints. Our goal with IHSMs is to eventually arrive at a system that, at low cost, can persist against a smart, well-funded adversary such as a secret service or organized cyber-crime. We apply Kerckhoff's principle and consider the attacker to have white-box access to the IHSM's hard-, firm- and software. We consider the attacker to have persistent access to the device and that they may be willing to spend weeks or months performing a single attack.

By targeting this pessimistic attacker model, we increase the real-world utility of IHSMs. Consider a group of healthcare providers intending to analyze a large database of patient health information. Accumulating potentially millions of sensitive medical records on a single system for such processing poses an inherent risk as this system becomes a valuable target

for organized cyber-criminals looking for ransom. IHSMs permit a level of physical security against e.g. a bribed insider that is as good as the level of network security afforded by modern firewalls and cryptographic protocols.

On the other end of the spectrum, consider a real-time group video communication provider. Relaying and transcoding video data between participants is hard to solve without trusting the server, but at the same time latency requires that the server is physically located close to its users. Given the global history of privacy-invasive cyber-attacks by secret services and other well-funded attackers, this may pose an issue. In this scenario, IHSMs enable the secure deployment of trusted server components closer to the user, or even at the network edge, where physical security is challenging.

An application with a similar scenario is manipulation-proof audit logging. Since IHSMs are connected to backup power, they can continue to record log messages from other nearby devices even during catastrophic disruption such as large-scale power outages. In this use case, the IHSM assumes two functions: That of a trusted, highly available data storage and that of a trusted timestamping service.

3.2 Inertial HSM motion

Against the background of these use cases, we will now elaborate on the four questions we formulated in the introduction to this section, starting with that on *type of motion*. There are several ways how we can approach motion. Periodic, aperiodic and continuous motion could serve the purpose. There is also linear motion as well as rotation. We can also vary the degree of electronic control in this motion.

The primary constraint on an IHSM's motion pattern is that it needs to be (almost) continuous to not expose any weak spots during instantaneous standstill of the HSM. Additionally, it has to stay within a confined space. For space efficiency, linear motion would have to be periodic, like that of a pendulum. Such periodic linear motion will have to quickly reverse direction at its apex so the device is not stationary long enough for this to become a weak spot.

In contrast to linear motion, rotation is space-efficient and can be continuous if the axis of rotation is inside the device. When the axis is fixed, rotation will expose a weak spot close to the axis where tangential velocity is low. Faster rotation can lessen the security impact of this fact at the expense of power consumption and mechanical stress, but it can never

eliminate it. More effective mitigations are additional tamper protection at the axis and having the HSM perform a compound rotation that has no fixed axis.

High speed gives rise to large centrifugal acceleration, which poses the engineering challenge of preventing rapid unscheduled disassembly of the device, but it also creates an obstacle to any attacker trying to manipulate the device in what we call a *swivel chair attack* (see Section 4.3). An attacker trying to follow the motion would have to rotate around the same axis. By choosing a suitable angular frequency we can prevent an attacker from following the device's motion since doing so would subject them to impractically large centrifugal forces. Essentially, this limits the approximate maximum size and mass of an attacker under an assumption on tolerable centrifugal force.

In this chapter, we focus on rotating IHSMs for simplicity of construction. For our initial research, we focus on systems with a fixed axis of rotation due to their simple construction but we do wish to note the challenge of hardening the shaft against tampering that any production device would have to tackle.

3.3 Tamper detection mesh construction

IHSMs do not eliminate the need for a security barrier. To prevent an attacker from physically destroying the moving part, tamper detection such as a mesh is still necessary. In this subsection, we will consider ways to realize this security barrier. In industry, mesh membranes are commonly used for tamper detection. Such membranes are deployed in systems for a variety of use cases ranging from low-security payment processing to high-security certificate management. From this, we can conclude that a properly implemented mesh *can* provide a practical level of security. In contrast to this industry focus, academic research has largely focused on ways to fabricate enclosures that embed characteristics of a PUF as a means of tamper detection [251, 118]. By using stochastic properties of the enclosure material to form a PUF, such academic designs leverage signal processing techniques to improve the system's security level by a significant margin.

In our research, we focus on security meshes as our IHSM's tamper sensors. The cost of advanced manufacturing techniques and special materials used in fine commercial meshes poses an obstacle to small-scale manufacturing and academic research. The foundation of an IHSM security is that

by moving the mesh, even a primitive, coarse mesh such as one made from a low-cost PCB becomes very hard to attack in practice. This allows us to use a simple construction using low-cost components. Additionally, the use of a mesh enables us to only spin the mesh itself and its monitoring circuit and keep the payload inside the mesh stationary for reduced design complexity. Tamper sensing systems such as RF fingerprinting that monitor the entire volume of the HSM instead of only a thin boundary layer would not allow for this degree of freedom in an IHSM. They would instead require the entire IHSM to spin including its payload, which would entail costly and complex systems for data and power transfer from the outside to the spinning payload.

3.4 Braking detection

The security mesh is a critical component in the IHSM's defense against physical attacks, but its monitoring is only one half of this defense. The other half consists of a reliable and sensitive braking detection system. This system must be able to quickly detect any slowdown of the IHSM's rotation. Ideally, a sufficiently sensitive sensor is able to measure any external force applied to the IHSM's rotor and should already trigger a response at the first signs of a manipulation attempt.

While the obvious choice to monitor rotation would be a magnetic or optical tachometer sensor attached to the IHSM's shaft, this would be a poor choice for our purposes since optical and magnetic sensors are susceptible to contact-less interference from outside. We could use feedback from the motor driver electronics to determine the speed. When using a BLDC motor, the driver electronics precisely know the rotor's position at all times. However, this approach might allow for attacks at the mechanical interface between the mesh and the motor's shaft. If an attacker can decouple the mesh from the motor e.g. by drilling, laser ablation, or electrical discharge machining (EDM) on the motor's shaft, the motor could keep spinning at its nominal frequency while the mesh is already standing still.

Instead of a stator-side sensor, a rotor-side inertial sensor such as an accelerometer or gyroscope placed inside the spinning mesh monitoring circuit would be a good component to serve as an IHSM's tamper sensor. A gyroscope would need to be placed close to the IHSM's shaft where centrifugal force is low, and would directly measure changes in angular velocity. An accelerometer could be placed anywhere on the rotor and would measure

centrifugal acceleration.

Modern, fully integrated MEMS accelerometers are very precise. By comparing acceleration measurements against a model of the device's mechanical motion, deviations can quickly be detected. This limits an attacker's ability to tamper with the device's motion. It may also allow remote monitoring of wear of the device's mechanical components such as bearings: MEMS accelerometers are fast enough to capture vibrations, which can be used as an early warning sign of failing mechanical components [141, 223, 36, 63].

In a spinning IHSM, an accelerometer mounted at a known radius with its axis pointing radially will measure centrifugal acceleration. Centrifugal acceleration rises linearly with radius, and with the square of frequency: $a = \omega^2 r$. For a given accelerometer and target speed of rotation, the accelerometer's location should be chosen to maximize dynamic range. A key point here is that for speeds between 500 and 1000 rpm, centrifugal acceleration already becomes very large at a radius of just a few cm. At 1000 rpm $\approx$ 17 Hz and at a 10 cm radius, centrifugal acceleration already is above 1000 $\frac{\mathrm{m}}{\mathrm{s}}$ or 100 g. Due to this large acceleration, the off-axis performance of the accelerometer has to be considered. Suitable high-g accelerometers for the large accelerations found on the circumference of an IHSM's rotor are mostly used in automotive applications.

To evaluate the feasibility of accelerometers as tamper sensors we can use a simple benchmark. Let us assume an IHSM spinning at 1000 rpm. To detect any attempt to brake it below 500 rpm, we have to detect a difference in acceleration of a factor of $\frac{\omega_2^2}{\omega_1^2} = 4$. Even without maximizing the accelerometer's dynamic range through optimal placement, any commercial MEMS accelerometer will suffice. Only to detect slow deceleration, the sensor's drift characteristics may have to be taken into account.

In Section 6 below, we conduct an empirical evaluation of a commercial automotive high-g MEMS accelerometer for braking detection in our prototype IHSM.

3.5 Mechanical layout

With our IHSM's components taken care of, what remains to be decided is how to put together these individual components into a complete device. A basic spinning HSM might look as shown in Figure 4.2. Visible are the axis of rotation, an accelerometer on the rotating part that is used to detect

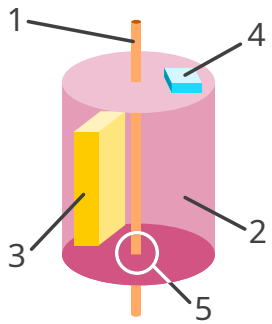

Figure 4.2: Concept of a simple spinning Inertial HSM. 1 - Shaft. 2 - Security mesh. 3 - Payload. 4 - Accelerometer. 5 - Shaft penetrating security mesh.

braking, the protected payload, and the area covered by the rotating tamper detection mesh. Note that we only have to move the tamper protection mesh, not the entire contents of the HSM, keeping most of the HSM's mass stationary. This reduces the moment of inertia of the rotating part. It also eliminates the need for rotating data and power connections to the payload, which can be supplied through a hollow shaft instead. In our proof-of-concept prototype, we accept a weak spot at the point where the shaft penetrates the mesh to simplify mechanical construction.

The spinning mesh must be designed to cover the entire surface of the payload, but it suffices if it sweeps over every part of the payload once per rotation. This means we can design longitudinal gaps into the mesh that allow outside air to flow through to the payload. In traditional boundary-sensing HSMs, cooling of the payload processor is a serious issue since any air duct or heat pipe would have to penetrate the HSM's security boundary. This problem can only be solved with complex and costly siphon-style constructions, so in commercial systems, heat conduction is used exclusively [123]. This limits the maximum power dissipation of the payload and thus its processing power. Using longitudinal gaps in the mesh, our setup allows direct air cooling of regular heatsinks. This unlocks much more powerful processing capabilities that greatly increase the maximum possible power dissipation of the payload. In an evolution of our design, the spinning mesh could even be designed to *be* a cooling fan.

Conventional HSMs are limited by the construction of their security meshes which rely on plastics as their main structural material. The security mesh has to fit the highest components inside the HSM. Since creating

a security mesh with a non-flat surface is difficult, this means there is an inevitable gap of a few millimeters between the surface of the payload CPU and the interior surface of the mesh. This distance is added to several millimeters of epoxy resin that the mesh must be embedded inside for it to be hard to remove intact. Overall, this leads to a structure approximately a centimeter thick that includes several millimeters of epoxy resin with particularly poor thermal conductivity [W194]. Even if "thermally conductive" resins would be used, thermal conductivity is limited to a fraction of what can be achieved with a heatsink directly attached to the CPU. A modern high-end CPU heatsink with its fan running has a thermal resistance from CPU junction to air of around 0.1 $\frac{\mathrm{K}}{\mathrm{W}}$ [W77]. If one were to make an HSM's security mesh out of an average thermally conductive epoxy with thermal conductivity $k \approx 1\,\frac{\mathrm{W\,K}}{\mathrm{m}}$ [142, 233, W171], the resulting thermal resistance for a 5-by-5 centimeter, 5 mm thermal interface alone would be 2 $\frac{\mathrm{K}}{\mathrm{W}}$, a more than 10-fold increase. For an acceptable temperature delta from junction to air of 60 K, this yields a maximum power dissipation of only 30 W compared to a theoretical 600 W for a conventional CPU cooler. Given that for modern high core-count CPUs both multithreaded performance and power dissipation are mostly linear in core count, this severely limits the achievable performance.

This estimated performance discrepancy matches up with our observation. Thales, a manufacturer of conventional HSMs reports 20 kOps/s ECC signature operations on NIST Curve P-256 on one of their top-of-range "Luna HSM 790" [W98], which compares to be slightly more than half of the 36 kOps/s signing operations that `openssl speed` in single-thread mode is able to do on an AMD Ryzen 7 PRO 4750U laptop CPU using 2.0 W of power on the active core. Using today's technology, we expect a performance jump of one to two orders of magnitude in computing power to be feasible in an IHSM compared to a conventional HSM.

3.6 Long-term Operation

Without settling on a particular design for an IHSM yet, from the previous sections we have already gained an understanding of what an IHSM would look like in practice. In the following paragraphs, we will draw some conclusions on how its design will affect the day-to-day operation of an IHSM. Like other HSMs, in a practical application, an IHSM may have to run continuously for a decade or even longer. As with any networked

system, a setup including IHSMs must be designed in a way that prevents the failure of one or several IHSMs on the network from compromising the whole system's security or reliability. Neither IHSMs nor traditional HSMs can withstand fire or flooding, so while a breach of security can be ruled out, a catastrophic failure of the device and erasure of data cannot [W107]. Traditionally, this problem is solved by storing all secrets in multiple, geographically redundant HSMs [W193]. On IHSMs this task is aided on the software layer since they are based on general-purpose computer hardware and allow for state-of-the-art database replication techniques to be applied without first porting them to an embedded operating system or foreign CPU architecture. A practical example of this approach is a 2019 technology demonstration [W159] created by signal.org, the organization running the signal secure messenger app. In this demonstration, signal.org have implemented the Raft consensus algorithm [199] inside Intel SGX to replicate state between geographically redundant enclaves.

Excluding natural disasters, there are three main categories of challenges to an IHSM's longevity: Failure of components of the IHSM due to age and wear, failure of the external power supply, and spurious triggering of the intrusion alarm by changes in the IHSM's environment. In the following paragraphs, we will evaluate each of these categories in their practical impact.

Component failure. The failure mode of an IHSM's components is the same as in any other computer system and the same generic mitigation techniques apply. The expected lifetime of electronic components can be increased by using higher-spec components and by reducing thermal, mechanical, and electrical stress. To reduce vibration stress on both rotor and stator, the rotor must be balanced. The main mechanical failure mode of an IHSM's is likely to be failure of the shaft bearings. By incorporating knowledge from other rotating devices that have a long lifetime such as cooling fans, this failure mode can be mitigated. Another noteworthy mechanical failure mode of an IHSM is dust buildup on the optical components of the communication link. This failure mode can be mitigated by routing cooling airflow such that it does not go past the communication link's optical components, as well as by filtering cooling air at the device's intakes.

Power failure. After engineering an IHSM's components to survive years of continuous operation, the next major failure mode to be considered is

power loss. Traditional HSMs solve the need for an always-on backup power supply by carrying large backup batteries [W194]. The low static power consumption of a traditional HSM's simple tamper detection circuitry allows for the use of non-replaceable backup batteries. An IHSM in contrast would likely require a rechargeable backup battery since its motor requires more power than the mesh monitoring circuit of a traditional HSM. In principle, a conventional Uninterruptible Power Supply (UPS) can be used, but in practice, a productized IHSM might have a smaller battery integrated. Conservatively assuming an average operating power consumption of 10 W for an IHSM's motor, a single large laptop battery with a capacity of 100 W h [W6] could already power an IHSM for 10 hours continuously. 10 W is a reasonable high estimate given that there are large industrial fans rated at lower wattages, e.g. Sunon CF2207LBL-000U-HB9, a 250 mm diameter 7.8 m^3/min axial fan rated at 6.6 W. If a built-in battery is undesirable or if power outages of more than a few seconds are unlikely (e.g. because of an external UPS), the IHSM's rotor itself can be used as a flywheel for energy storage.

Spurious alarms due to vibration. Even with all components working to their specification, an IHSM could still catastrophically fail if for some reason its alarm would be spuriously activated due to movement of the device. The likelihood of such an alarm failure must be minimized, e.g. by employing vibration damping. There are several possible causes why an IHSM might move during normal operation. The IHSM may have to be relocated between data centers, or a worker may bump the IHSM. Additionally, the effect of normal mechanical vibration on the IHSM's tamper sensors has to be considered. During normal operation, vibration from outside sources such as backup generators and nearby traffic (e.g. trains) may couple into the IHSM through the building. Since IHSMs are rotating machines they will themselves cause some amount of vibration and thus vibration isolation is a reasonable design requirement. Besides everyday sources of mechanical noise, (usually harmless) earthquakes are a common occurrence in some regions of the world and will couple through any reasonable amount of vibration damping.

None of these sources of mechanical noise are likely to cause a false alarm. For reference, consider an IHSM running at an angular velocity of 1000 rpm. A tamper sensor mounted at a radius of 100 mm will measure a constant centrifugal acceleration of approximately 100 g. Literature on car

crashes shows that accelerations above $10\,g$ in the car's structural components correspond to a crash at $30\,\frac{\mathrm{km}}{\mathrm{h}}$ and above [W3, 88]. Measurements of the Peak Ground Acceleration (PGA) of severe earthquakes show that even the strongest earthquakes rarely reach a PGA of 0.1 g [76] with the 2011 Tohoku earthquake at approximately 0.3 g.

Instantaneous acceleration increases linearly with frequency, but likewise, simple vibration dampers work better with higher frequencies [131, 22, 56], To reduce the likelihood of false detections, it is enough to damp high-frequency shock and vibration, as low-frequency shock or vibration components will not reach accelerations large enough to cause a false alarm. For instance, an earthquake's low-frequency vibrations dissipate a tremendous amount of mechanical power across a large geographic area, but due to their low absolute instantaneous acceleration, we can ignore them for the purposes of our tamper detection system. An IHSM's tamper detection subsystem will be able to clearly distinguish attempts to stop the IHSM's rotation from normal environmental noise by their magnitude. Any external acceleration that would come close in order of magnitude to the operating centrifugal acceleration at the periphery of an IHSM's rotor would likely destroy the IHSM.

3.7 Transportation

While unintentional acceleration is unlikely to cause false alarms in an IHSM when simple vibration damping is employed, there is an issue when intentionally moving an IHSM: The IHSM's rotor stores significant rotational energy and will respond to tipping with a precession force. This could become an issue when a larger IHSM is transported between e.g. the manufacturer's premises and its destination data center. The simple solution to this problem is to transport the IHSM elastically mounted with its axis pointing upwards inside a shipping box that is weighted to resist precession forces.

During shipping, the IHSM will require a continuous power supply. Following our conservative estimate in Section 3.6, 48-hour courier shipping could easily be bridged with the equivalent of 5-10 laptop batteries. In applications that do not require a backup battery built-in to the IHSM (e.g. due to existing UPS backup), the IHSM could be shipped connected to an external battery akin to a "power bank" that is sent back to the IHSM's manufacturer after the IHSM has been installed. Long-distance shipping

can be facilitated through compatibility with standards used for powered refrigerated shipping containers.

3.8 Graceful Failover and Maintenance

As described above, failure can never be fully prevented. However, finely-grained monitoring of operational parameters may be capable of recognizing some types of failure such as backup battery failure, mechanical wear, or over/undertemperature conditions some time before alarm levels have been reached and all secrets must be destroyed. This type of early warning allows for the implementation of a graceful failover mechanism. Similar to hot spares in hard disk arrays, a number of IHSMs might share a hot spare IHSM that is running, but that does not yet contain any secrets. Once an IHSM detects early warning signs of an impending failure, it can then transfer its secrets to the hot spare using replication technologies as mentioned in the previous paragraph, then delete its local copies. This would allow for the graceful handling of device failures due to both age and disasters such as fires.

When such failovers happen, IHSMs provide a key benefit compared to traditional HSMs. Since an IHSM is not permanently potted and its security mesh is mechanically robust, it can be stopped and disassembled to repair a faulty component such as a worn-out bearing or a defective payload component such as a RAM module or an SSD. A faulty IHSM can be refurbished like a normal server. Its disassembly does not require any special equipment.

The primary challenge in repairing IHSMs is purely operational. It has to be ensured that an attacker lying in wait cannot seize the opportunity of the IHSM's defenses shutting down to implant a hardware trojan. A possible approach would be to have the IHSM contain a cryptographic identity that it uses to authenticate its status to its operator, and that is destroyed along with the payload's secrets when the IHSM is tampered with. The IHSM's operator could then provide a cryptographically authenticated maintenance token to a trusted technician that allows the technician to power down this particular IHSM during a set time window. The technician can then physically repair the IHSM and return it into service, after which the operator can use the IHSM's identity to verify that the repair was conducted as intended. Using a physical token instead of powering off the IHSM remotely prevents the accidental unsupervised stopping of an IHSM due to operator

error.

To decrease the risk posed by a rogue technician, similar to the DNSSEC root key signing ceremonies [W221], arbitrarily complex procedures can be implemented that could, for example, require each maintenance procedure to be accompanied by several independent witnesses.

4 Attacks

After outlining the basic mechanical design of an inertial HSM as well as the fundamentals of its long-term operation above, in this section, we will detail possible ways to attack it. At the core of an IHSM's defenses is the same security mesh or other technology as it is used in traditional HSMs. This means that ultimately an attacker will have to perform the same steps they would have to perform to attack a traditional HSM. However, they will either need to perform these attack steps with a tool such as a CNC actuator or a laser that follows the HSM's rotation at high speed, or they will first need to defeat the braking sensor.

4.1 Attacks that don't work

In the sections below, we will go into detail on such attacks on IHSMs. To put these attack approaches into perspective, we will start with a brief overview of attacks on conventional HSMs that the IHSM is defended against.

In principle, there are three ways to attack a conventional HSM. The hard way is to go through the security mesh without triggering the alarm, e.g. with a probe that is finer than the mesh's spacing. For larger probes, an attacker can laboriously uncover, then bridge the mesh traces to allow part of the mesh to be removed. Some HSMs attempt to detect such attacks by measuring mesh resistance [W194], but this is limited by available measurement precision. If an attacker only wishes to disable a small section of the mesh to insert a handful of fine probes into the device, this hardening approach becomes challenging. Consider a mesh that covers an area of 100 mm by 100 mm. An attacker who short-circuits a 5 mm by 5 mm section of this mesh will change the mesh trace's resistance by approximately $\frac{5\,\text{mm}\cdot 5\,\text{mm}}{100\,\text{mm}\cdot 100\,\text{mm}} = 0.25\%$. Detecting this change would require a resistance measurement of at least 9 bit of precision and corresponding temperature stability of the mesh material.

The second way to attack an HSM is to go *around* the mesh. Many

commercial HSMs sandwich the payload PCB between two halves of an enclosure [W194]. This design is vulnerable to attempts to stick a fine needle through the interface between lid and PCB [W189]. Conventional HSMs mitigate this weak spot by wrapping a patterned conductive foil around the HSM that forms the security mesh, leaving only the corners and the payload's power and data feed-through as potential weak spots.

The third and last way to attack a conventional HSM is to disable the mesh monitoring circuit [W189]. An attacker may need to insert several probes to wiretap the payload processor's secrets, but if poorly implemented, they may be able to disable the mesh monitor with only one. This type of attack can be mitigated by careful electronic design that avoids single points of failure as well as fail-open failure modes.

4.2 Attacks that work on any HSM

An IHSM provides an effective mitigation against direct attacks on the security mesh as described in the previous paragraphs. However, there are certain generic attacks that work against any HSM technology, conventional or IHSM. One type of these attacks are contactless attacks such as electromagnetic (EM) side-channel attacks. EM side-channel attacks can be mitigated by shielding and by designing the IHSM's payload such that critical components such as CPUs are physically distant to the security mesh, preventing EM probes from being brought close. Conducted EMI side-channels that could be used for power analysis can be mitigated by placing filters on the inside of the security mesh at the point where the power and network connections penetrate the mesh [14]. Finally, the API between the HSM's payload and the outside world provides attack surface. Attacks through the network interface must be prevented as in any other networked system by only exposing the minimum necessary amount of API surface to the outside world, and by carefully vetting this remaining attack surface [14].

IHSMs do not provide an inherent benefit against such contactless attacks. However, there are two mitigating factors in play that still give IHSMs an advantage over conventional HSMs in this scenario. Because IHSM meshes can be made using simpler technology than conventional HSM meshes at the same level of security, IHSMs can use larger meshes and are less space-constrained. This larger volume allows for a greater physical distance between security-critical components and places accessible to an attacker using an electromagnetic probe for EM side-channel attacks.

Another type of attack that is possible against all types of HSMs are software attacks. Flaws in an HSM's software such as memory safety errors in its external-facing APIs can lead to a full compromise of the HSM's secrets [24]. Like a traditional HSM, an IHSM has to expose some API to the outside world to be useful. For both, the hardening techniques are the same as in any other networked system and include the reduction of attack surface e.g. through firewalling, fuzz testing, and formal verification. In IHSMs these mitigations are easier to implement since they allow the use of conventional server hardware and well-audited open source software, instead of hard-to-audit proprietary code on an embedded platform.

4.3 The Swivel Chair Attack

If we assume whoever integrates the payload into an IHSM has done adequate work and prevented all contactless attacks, we are left with attacks that aim at mechanically bypassing the IHSM's security mesh. The first type of attack we will consider is the most basic of all attacks: a human attacker holding a soldering iron trying to rotate herself along with the mesh using a very fast swivel chair. Let us pessimistically assume that this co-rotating attacker has their center of mass on the axis of rotation. The attacker's body is likely on the order of 200 mm wide along its shortest axis, resulting in a minimum radius from axis of rotation to surface of about 100 mm. Wikipedia lists horizontal g forces in the order of 20 g as the upper end of the range tolerable by humans for a duration of seconds or above. We thus set our target acceleration to $100\,\mathrm{g} \approx 1000\,\mathrm{m/s^2}$, a safety factor of 5 past that range. Centrifugal acceleration is $a = \omega^2 r$. In our example, this results in a minimum angular velocity of $f_{\min} = \frac{1}{2\pi}\sqrt{\frac{a}{r}} = \frac{1}{2\pi}\sqrt{\frac{1000\,\mathrm{m/s^2}}{100\,\mathrm{mm}}} \approx 16\,\mathrm{Hz} \approx 1000\,\mathrm{rpm}$. From this, we can conclude that even at moderate speeds of 1000 rpm and above, a manual attack is no longer possible and any attack would have to be carried out using some kind of mechanical tool. Literature supports this conclusion, with loss of orientation reported as early as at 70 rpm in an observer located on the axis of rotation [71].

Figure 4.3 shows a schematic overview of the structure of such a rotating attack tool. The tool itself has to rotate at the IHSM's speed because counter-rotating the IHSM instead, the accelerometer on the rotor would measure lower centrifugal acceleration and detect the manipulation attempt. Following the IHSM's rotation closely enough to allow for remote-controlled

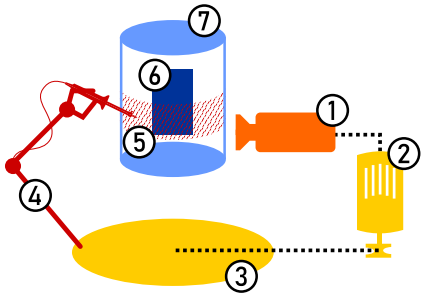

Figure 4.3: Schematic overview of a robotic rotating-stage attack. An optical sensor (1) observes the IHSM's rotation and adjusts the setpoint of a servo motor (2) that rotates the attack stage (3). On the rotating attack stage, a remote-controlled manipulator (4) is mounted that deactivates the security mesh (7) and creates an opening (5). Through this opening, a human operator can then insert tools such as probes to read out sensitive information from the actual payload (6).

manipulation of the IHSM is hard. Let us assume a small IHSM mesh with radius $r = 100\,\text{mm}$. To keep a manipulator stationary within a 5 mm by 5 mm window over a period of 10 s requires attack tool and IHSM speeds to be matched to an accuracy better than $\frac{5\,\text{mm}}{10\,\text{s}} \cdot \frac{1}{2\pi r} = 8.0\,\text{mHz} = 0.048\,\text{rpm}$. Relative to a realistic IHSM's speed of 1000 rpm this corresponds to approximately 50 ppm. Achieving this accuracy would likely require active servo control of the attack tool's rotation that is locked, e.g. optically, to the IHSM's rotor.

If an attacker were to solve the tracking issue, the remaining issue is that they still need to construct a remote-controlled manipulator that is able to disable the IHSM's mesh. This manipulator would have to be tolerant to high g forces so that it can be mounted on the attack tool's rotating stage. Drilling only a small hole is not enough in this case since, while the mesh is moving, the payload is stationary. Instead, using the rotating manipulator, the attacker has to create an opening in the mesh large enough to place a *stationary* probe on the payload. We estimate that creating a rotating, remote-controllable manipulator that can be used to successfully attack a security mesh is infeasible given the degree of manual skill necessary even for normal soldering work.

4.4 Mechanical weak spots

As we elaborated in the previous paragraphs, we consider a fast-moving mesh to offer a strong tamper detection capability based on the assumption

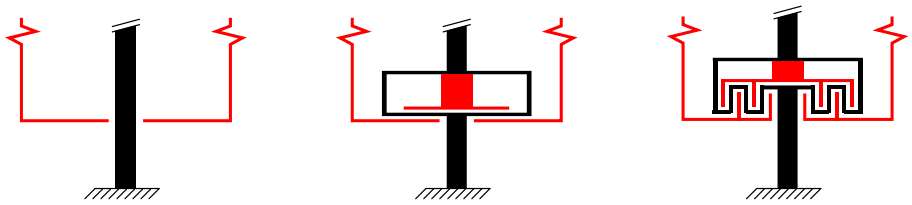

(a) Cross-sectional view of the basic configuration with no special protection of the shaft. Red: moving mesh – Black: stationary part.

(b) An internal, independently rotating disc greatly decreases the space available to attackers at the expense of another moving part and a second moving monitoring circuit.

(c) A second moving tamper detection mesh also enables more complex topographies.

Figure 4.4: Mechanical countermeasures to attacks through or close to the shaft of a fixed-axis rotating IHSM.

that the mesh is moving too fast to tamper. However, depending on the type of motion used, the mesh's actual speed may vary by location and over time. Our example configuration of a rotating mesh moves continuously and does not have any time-dependent weak spots. It does, however, have a weak spot where the shaft penetrates the mesh at the axis. The mesh's tangential velocity decreases close to the shaft, and the shaft itself may allow an attacker to insert tools such as probes into the device through the opening it creates. Conventional HSMs also have to take precautions to protect their power and data connections. In conventional HSMs, power and data are routed into the enclosure along a meandering path through the PCB or through flat flex cables sandwiched in between security mesh foil layers [236]. As a result of these precautions, in conventional HSMs, this interface rarely is a mechanical weak spot. In inertial HSMs, careful engineering is necessary to achieve the same effect. Figure 4.4 shows variations of the shaft interface with increasing complexity.

4.5 Attacking the mesh in motion

To disable the mesh itself, an attacker can choose two paths. One is to attack the mesh itself, for example by bridging its traces. The other option is to tamper with the monitoring circuit to prevent a damaged mesh from triggering an alarm [W189]. Attacks in both locations require electrical contact to parts of the circuit. Traditionally, this is done by soldering a wire or by placing a probe. We consider this type of attack hard to perform

on an object spinning at high speed. Possible remaining attack avenues may be to rotate an attack tool in sync with the mesh or to use a laser or ion beam fired at the mesh to cut traces or carbonize parts of the substrate to create electrical connections. Encapsulating the mesh in a potting compound and shielding it with a metal enclosure as is common in traditional HSMs will significantly increase the complexity of such attacks.

4.6 Attacks on the rotation sensor

Instead of attacking the mesh in motion, an attacker may also try to first stop the rotor. To succeed, they would need to falsify the rotor's MEMS accelerometer measurements. We can disregard electronic attacks on the sensor or the monitoring microcontroller because they would be no easier than attacking the mesh traces. What remains would be physical attacks of the accelerometer's sensing mechanism. In a MEMS accelerometer, a proof mass moves a cantilever whose precise position is measured electronically. A topic of recent academic interest has been acoustic attacks tampering with these mechanics [254], but such attacks do not yield sufficient control to precisely falsify sensor readings. A possible more invasive attack may be to first decapsulate the sensor MEMS using laser ablation synchronized with the device's rotation. Then, a fast-setting glue such as a cyanoacrylate could be deposited on the MEMS, locking the mechanism in place. This type of attack can be mitigated by mounting the accelerometer in a shielded location inside the security envelope and by varying the rate of rotation over time.

4.7 Attacks on the alarm circuit

Besides trying to deactivate the tamper detection mesh, an electronic attack could also target the alarm circuitry inside the stationary payload or the communication link between rotor and payload. The link can be secured using a cryptographically secured protocol like one would use for wireless radio links along with a high-frequency heartbeat message. The alarm circuitry has to be designed such that it is entirely contained within the HSM's security envelope. Like in conventional HSMs, it has to be built to either tolerate or detect environmental attacks using sensors for temperature, ionizing radiation, laser radiation, supply voltage variations, ultrasound or other vibration, and gases or liquids. If a wireless link is used between the IHSM's rotor and stator, this link must be cryptographically secured. To

prevent replay attacks, link latency must continuously be measured, so this link must be bidirectional.

4.8 Fast and violent attacks

A variation of the above attacks on the alarm circuitry is to use a tool such as a large hammer or a gun to simply destroy the part of the HSM that erases data in response to tampering before it can perform its job. To mitigate this type of attack, the HSM must be engineered to be either tough or brittle: Tough enough that the tamper response circuitry will reliably withstand any attack for long enough to carry out its function or brittle in a way that during any attack, the payload is reliably destroyed before the tamper response circuitry.

5 Proof-of-concept Prototype implementation

As we elaborated above, the mechanical component of an IHSM significantly increases the complexity of any attack even when implemented using only common, off-the-shelf parts. In view of this amplification of design security, we have decided to validate our theoretical studies by implementing a proof-of-concept prototype IHSM (Figure 4.1). The main engineering challenges we set out to solve in this proof-of-concept prototype were:

1. A mechanical design suitable for rapid prototyping that can withstand at least 500 rpm.
2. The automatic generation of security mesh PCB layouts for quick adaption to new form factors.
3. Non-contact power transmission from stator to rotor.
4. Non-contact bidirectional data communication between stator and rotor.

We will outline our findings on these challenges one by one in the following paragraphs.

5.1 Mechanical design

We sized our proof-of-concept prototype to have sufficient payload space for a Raspberry Pi single-board computer to approximate a traditional HSM's

processing capabilities. We use printed circuit boards as the main structural material for the rotating part, and 2020 aluminium extrusion for its mounting frame. Figure 4.5 shows the rotor's mechanical PCB designs. The design uses a 6 mm brass tube as its shaft, which is sufficiently narrow to pose a challenge to an attacker. The rotor is driven by a small hobby quadcopter motor. Our prototype incorporates a functional PCB security mesh. As we observed previously, this mesh only needs to cover every part of the system once per revolution, so we designed the longitudinal PCBs as narrow strips to save weight.

5.2 PCB security mesh generation

Our proof-of-concept security mesh covers a total of five interlocking mesh PCBs (Figure 4.6b). A sixth PCB contains the monitoring circuit and connects to these mesh PCBs. To speed up design iterations, we automated the generation of this security mesh through a plugin for the KiCAD EDA suite[1]. Figure 4.6a visualizes the mesh generation process. First, the target area is overlaid with a grid. Then, the algorithm produces a randomized tree covering the grid. Finally, individual mesh traces are traced according to a depth-first search through this tree. We consider the quality of the plugin's output sufficient for practical applications. Together with FreeCAD's KiCAD StepUp plugin, this results in an efficient toolchain from mechanical CAD design to production-ready PCB files.

5.3 Power transmission from stator to rotor

The spinning mesh has its own autonomous monitoring circuit. This spinning monitoring circuit needs both power and data connectivity to the stator. To design the power link, we first need to estimate the monitoring circuit's power consumption. We base our calculation on the (conservative) assumption that the spinning mesh sensor should send its tamper status to the static monitoring circuit at least once every $T_{\mathrm{tx}} = 10\,\mathrm{ms}$. At 100 kBd, a transmission of a one-byte message in standard UART framing would take 100 µs and yield a 1 % duty cycle. If we assume an optical or RF transmitter that requires 10 mA of active current, this yields an average operating current of 100 µA. This value is comparable to a reasonable estimation of the current consumption of the monitoring circuit itself. In our prototype, we used ST Microelectronics STM32 Series ARM

[1] https://blog.jaseg.de/posts/kicad-mesh-plugin/

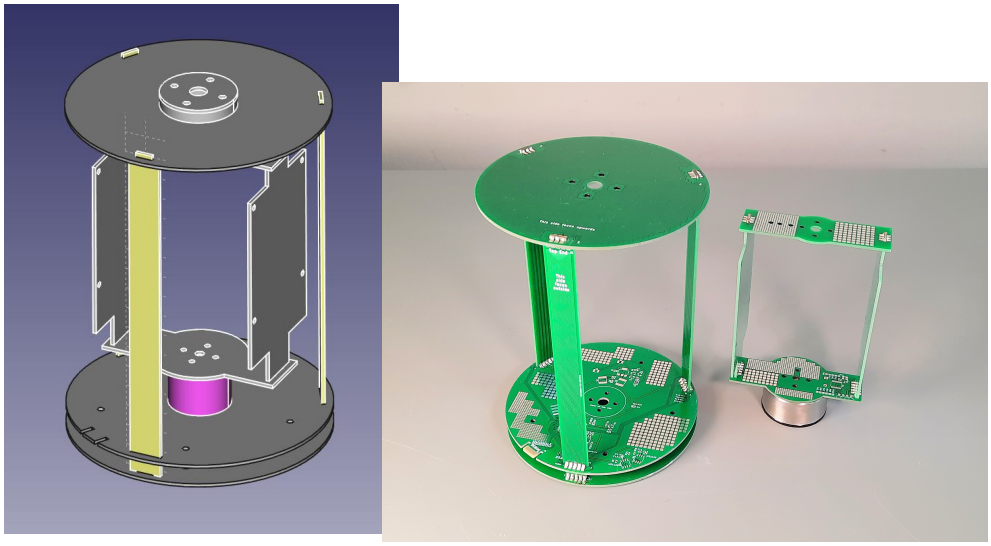

(a) The 3D CAD design of the proof-of-concept prototype.

(b) Assembled mechanical prototype rotor (left) and stator (right) PCB components.

Figure 4.5: Our proof-of-concept prototype IHSM's PCB security mesh design

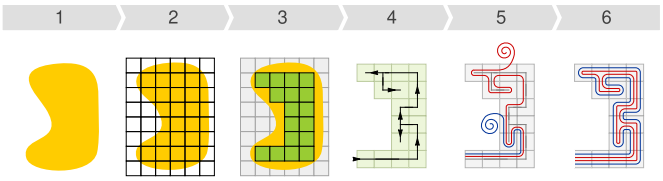

(a) Overview of the automatic security mesh generation process. 1 - Example target area. 2 - Grid overlay. 3 - Grid cells outside of the target area are removed. 4 - A random tree covering the remaining cells is generated. 5 - The mesh traces are traced along a depth-first walk of the tree. 6 - Result.

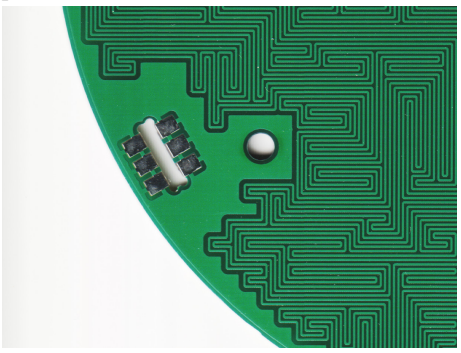

(b) Detail of a PCB produced with a generated mesh.

Figure 4.6: Our automatic security mesh generation process

Cortex-M microcontrollers. To get an estimate on the current consumption of an energy-optimized design we will refer to the datasheet of the STM32L486JG[2], a representative member of ST's STM32L4 low-power subfamily that provides hardware acceleration for AES256. A good target for an implementation of a secure cryptographic channel on this device would be the noise protocol framework [W205]. While the initial handshake for key establishment uses elliptic-curve cryptography and may take several hundred milliseconds [258], the following payload data transfer messages require only symmetric cryptographic primitives. The STM32L486JG datasheet lists the microcontroller's typical operating current at around 8 mA at 48 MHz clock speed and lists a sleep current of less than 1 µA in low-power standby mode with RTC enabled. The AES peripheral is listed with less than 2 $\frac{\mu A}{MHz}$ typical current consumption. A typical high-g accelerometer for an IHSM application would be ST Microelectronics' H3LIS331DL. Its datasheet[3] lists a typical current consumption between 10 µA in low-power mode with output sampling rate up to 10 Hz and 300 µA in normal operating mode with output sampling rate up to 1 kHz. Given the low amount of data that has to be processed in our application (hundreds of bytes per second), if we assume a duty cycle of 1 % of active data processing vs. sleep mode at the given clock speed we arrive at a total estimated current consumption of our microcontroller of less than 100 µA. Thus, reserving 100 µA for the monitoring circuit on top of the 100 µA for the transceiver circuit we arrive at an energy consumption of 1.7 A h per year.

This annual energy consumption is close to the capacity of a single CR123A lithium primary cell. By either using several such cells or by optimizing power consumption, several years of battery life could easily be reached. In our proof of concept prototype, we decided against using a battery to reduce rotor mass and avoid balancing issues.

We also decided against mechanically complex solutions such as slip rings or electronically complex ones such as inductive power transfer. Instead, we chose a simple setup consisting of a stationary lamp pointing at several solar cells on the rotor. At the monitoring circuit's low power consumption power transfer efficiency is irrelevant, so this solution is practical. Our system uses six series-connected solar cells mounted on the end of the cylindrical rotor that are fed into a large 33 µF ceramic buffer capacitor through a Schottky diode. This solution provides around 3.0 V at several

[2] https://www.st.com/resource/en/datasheet/stm32l486jg.pdf
[3] https://www.st.com/resource/en/datasheet/h3lis331dl.pdf

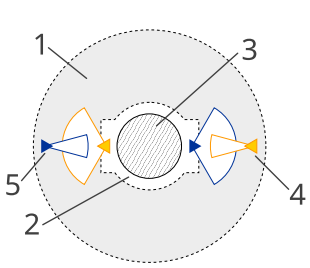

(a) Basic layout, view along axis of rotation. 1 - Rotor base PCB. 2 - Stator IR link PCB. 3 - Motor. 4 - receiver PIN photodiode. 5 - transmitter IR LED.

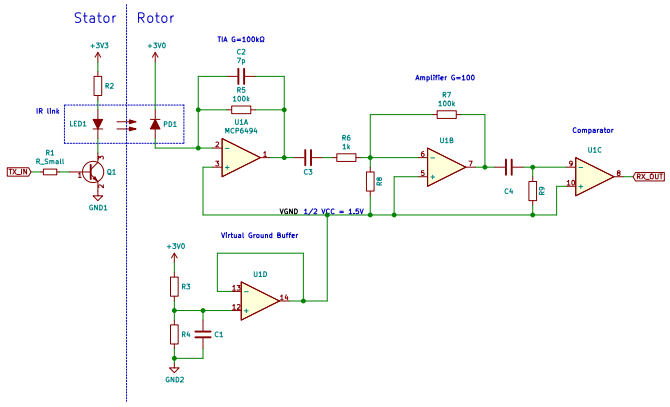

(b) Schematic with sample component values. C2 is highly dependent on the photodiode characteristics and stray capacitances.

Figure 4.7: IR data link implementation

tens of mA to the payload when illuminated using either a 60 W incandescent light bulb or a flicker-free LED studio light of similar brightness[4].

5.4 Data transmission between stator and rotor

Besides power transfer from stator to rotor, we need a reliable, bidirectional data link to transmit mesh status and a low-latency heartbeat signal. We chose to transport an 115 kBd UART signal through a simple IR link for a quick and robust solution. The link's transmitter directly drives a standard narrow viewing angle IR led through a transistor. The receiver has an IR PIN photodiode reverse-biased at $\frac{1}{2}V_{CC}$ feeding into an MCP6494 general purpose opamp configured as a 100 kΩ transimpedance amplifier. As shown in Figure 4.7b, the output of this TIA is amplified one more time before being squared up by a comparator. Our design trades off stator-side power consumption for a reduction in rotor-side power consumption by using a narrow-angle IR led and photodiode on the rotor, and wide-angle components at a higher LED current on the stator. Figure 4.7a shows the physical arrangement of both links. The links face opposite one another and are shielded from one another by the motor's body in the center of the PCB.

[4]LED lights intended for room lighting exhibit significant flicker that can cause the monitoring circuit to reset. Incandescent lighting requires some care in shielding the data link from the light bulb's considerable infrared output.

5.5 Evaluation

The proof-of-concept hardware worked as intended. Both rotating power and data links performed well. As we expected, the mechanical design vibrated at higher speeds but despite these unintended vibrations, we were able to reach speeds in excess of 1000 rpm by clamping the device to the workbench. Even at high speeds, both the power link and the data links continued to function without issue.

By design, our prototype is not yet a production-ready solution. Its main limitation is the small payload volume that can house one or two Raspberry Pi single-board computers but does not allow for more powerful hardware such as a contemporary server mainboard. Being constructed without access to a proper mechanical workshop, its imprecise construction leads to vibration at high speeds. Its optical communication links in breadboard construction function and need to be translated into manufacturable PCBs, and its security mesh has to be optimized for security. Finally, a motor driver solution needs to be selected that allows for direct digital control of motor speed. Overall, the prototype soundly demonstrated the viability of the IHSM concept and we are confident that all of these limitations can be conclusively solved in a new iteration that might be a "beta" version of a practical IHSM, built in a mechanical workshop.

6 Using MEMS accelerometers for braking detection

In our proof-of-concept prototype, for braking detection we chose an accelerometer placed on the circumference of our prototype's rotor for two reasons: First, it avoids the likely issue of high centrifugal acceleration falsifying gyroscope measurements. Second, by orienting one axis of the accelerometer radially, we can avoid exceeding the accelerometer's range even when rapidly accelerating or decelerating. Rapid angular acceleration or deceleration produces high tangential linear acceleration or deceleration in our sensor, but the radially-oriented axis of the accelerometer only experiences an amount of centrifugal acceleration that is bounded by the rotor's momentary angular velocity and never exceeds the device's specified operating conditions.

Using our prototype, we performed an evaluation of an AIS1120 commercial automotive MEMS accelerometer as a braking sensor. The device

is mounted inside our prototype at a radius of 55 mm from the axis of rotation to the center of the device's package. The AIS1120 provides a measurement range of $\pm 120\,g$. At its 14-bit resolution, one LSB corresponds to $15\,\mathrm{m}g$.

Our prototype IHSM uses a motor controller intended for use in RC quadcopters. In our experimental setup, we manually control this motor controller through an RC servo tester. In our experiments, we externally measured the device's speed of rotation using a magnet fixed to the rotor and a reed switch. The reed switch output is digitized using a USB logic analyzer at a sample rate of 100 MHz. We calculate rotation frequency as a 1 s running average over interval lengths of the debounced captured signal[5].

The accelerometer is controlled from the STM32 microcontroller on the rotor of our IHSM prototype platform. Timed by an external quartz, the microcontroller samples accelerometer readings at 10 Hz. Readings are accumulated in a small memory buffer, which is continuously transmitted out through the prototype platform's infrared link. Data is packetized with a sequence number indicating the buffer's position in the data stream and a CRC-32 checksum for error detection. On the host, a Python script stores all packets received with a valid checksum in an SQLite database.

Data analysis is done separately from data capture. An analysis IPython notebook reads captured packets and reassembles the continuous sample stream based on the packets' sequence numbers. The low 10 Hz sample rate and high 115 kBd transmission speed lead to a large degree of redundancy with gaps in the data stream being rare. This allowed us to avoid writing retransmission logic or data interpolation.

Figure 4.8a shows an entire run of the experiment. During this run, we started with the rotor at standstill, then manually increased its speed of rotation in steps. Areas shaded gray are intervals where we manually adjust the rotor's speed. The unshaded areas in between are intervals when the rotor speed is steady. Figure 4.8b shows a magnified view of these periods of steady rotor speed. In both graphs, orange lines indicate centrifugal acceleration as calculated from rotor speed measurements. Visually, we can see that measurements and theory closely match. Our frequency measurements are accurate and the main source of error are the accelerometer's intrinsic errors as well as error in its placement due to construction tolerances.

The accelerometer has two main intrinsic errors. Offset error is a fixed

[5] A regular frequency counter or commercial tachometer would have been easier, but neither was available in our limited COVID-19 home office lab.

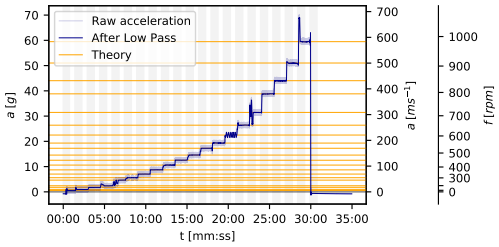

(a) Raw recording of accelerometer measurements during one experiment run. Shaded areas indicate time intervals when we manually adjusted speed.

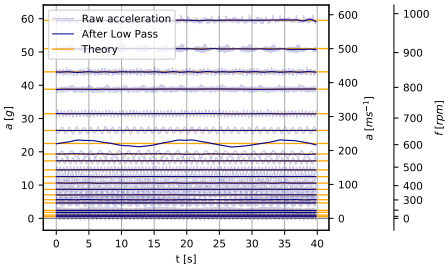

(b) Valid measurements cropped out from 4.8a for various frequencies. Intermodulation artifacts from the accelerometer's 10 Hz sampling frequency and the 3 Hz to 18 Hz rotation frequency due to gravity and device vibration are clearly visible.

Figure 4.8: Traces of acceleration measurements during one experiment run.

additive offset to all measurements. Scale error is an error proportional to a measurements value that results from a deviation between the device's specified and actual sensitivity. We correct for both errors by first extracting all stable intervals from the time series, then fitting a linear function to the measured data. Offset error is this linear function's intercept, and scale error is its slope. We then apply this correction to all captured data before plotting and later analysis. Despite its simplicity, this approach already leads to a good match of measurements and theory modulo a small part of the device's offset remaining. At high speeds of rotation, this remaining offset does not have an appreciable impact, but due to the quadratic nature of centrifugal acceleration, at low speed, it causes a large relative error of up to 8 % at 95 rpm.

After offset and scale correction, we applied a low-pass filter to our data. The graphs show both raw and filtered data. Raw data contains significant harmonic content. This content is due to vibrations in our prototype as well as gravity since we tested our proof-of-concept prototype lying down, with its shaft pointing sideways. FFT analysis shows that this harmonic content is a clean intermodulation product of the accelerometer's sample rate and the speed of rotation with no other visible artifacts.

Figure 4.9 shows a plot of our measurement results against frequency. Data points are shown in dark blue, and theoretical behavior is shown in orange. From our measurements, we can conclude that an accelerometer is a good choice for an IHSM's braking sensor. A simple threshold set according to the sensor's calculated expected centrifugal force should be sufficient to

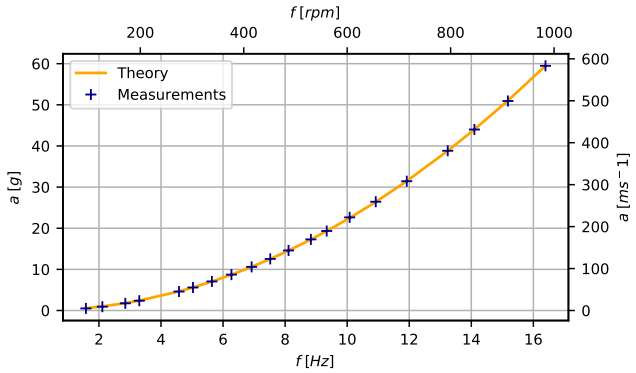

Figure 4.9: Centrifugal acceleration versus angular frequency in theory and in our experiments. Experimental measurements are shown after correction for offset and scale error. Above 300 rpm, the relative error is below 0.5 %. Below 300 rpm, the residual offset error has a large impact (0.05 g absolute or 8% relative at 95 rpm.)

reliably detect manipulation attempts without resulting in false positives. Periodic controlled changes in the IHSM's speed of rotation allow offset and scale calibration of the accelerometer on the fly, without stopping the rotor.

7 Conclusion

In this chapter, we introduced Inertial Hardware Security Modules (IHSMs), a novel concept for the construction of advanced hardware security modules from simple components. We analyzed the concept for its security properties and highlighted its ability to significantly strengthen otherwise weak tamper detection barriers. We validated our design by creating a proof-of-concept hardware prototype. In this prototype, we have demonstrated practical solutions to the major electronics design challenges: Data and power transfer through a rotating joint, and mechanized mesh generation. We have used our prototype to perform several experiments to validate the rotary power and data links and the onboard accelerometer. Our measurements have shown that our proof-of-concept solar cell power link works well and that our simple IR data link already is sufficiently reliable for telemetry. Our experiments with an AIS1120 automotive MEMS accelerometer showed that this part is well-suited for braking detection in the range of rotation speed relevant to the IHSM scenario.

Overall, our findings validate the viability of IHSMs as an evolutionary step beyond traditional HSM technology. IHSMs offer a high level of se-

curity beyond what traditional techniques can offer even when built from simple components. They allow the construction of devices secure against a wide range of practical attacks in small quantities and without specialized tools. The rotating mesh allows longitudinal gaps, which enables new applications that are impossible with traditional HSMs. Such gaps can be used to integrate a fan for air cooling into the HSM, allowing the use of powerful computing hardware inside the HSM. We hope that this simple construction will stimulate academic research into (more) secure hardware.

Building on the foundations of IHSM construction that we laid out in this chapter, in the following two chapters we will provide detailed solutions for two key design challenges in IHSM construction. In Chapter 5, we will introduce a low-cost tamper sensing mesh monitoring circuit based on Time Domain Reflectometry. Using this approach, we can further strengthen the security of meshes created using simple manufacturing processes in an IHSM. In Chapter 6, we approach the question of a rotation-invariant wireless inductive power supply for an IHSM and provide a planar inductor layout that minimizes voltage ripple with IHSM rotation.

Web sources

[W3] *A Test Procedure for Airbags*. 2002 (cit. on p. 89).

[W6] US Federal Aviation Administration. *Pack Safe: Batteries, Lithium*. 2018-05-31. URL: https://www.faa.gov/hazmat/packsafe/more_info/?hazmat=7 (visited on 2021-07-12) (cit. on p. 88).

[W77] Emmanouil D. Fylladitakis. *Top Tier CPU Air Coolers Q3 2015: 9-Way Roundup Review*. URL: https://www.anandtech.com/show/9415/top-tier-cpu-air-coolers-9way-roundup-review/12 (visited on 2021-07-08) (cit. on p. 86).

[W98] Thales Group. *Thales Luna HSM Product Family Overview Page*. 2021. URL: https://cpl.thalesgroup.com/encryption/hardware-security-modules/network-hsms (visited on 2021-07-08) (cit. on p. 86).

[W100] Lester Haines. *US Outfit Patents 'invisible' UAV: Stealth through Persistence of Vision*. Ed. by The Register. 2006-09-25. URL: https://www.theregister.com/2006/09/25/phantom_sentinel/ (visited on 2020-09-17) (cit. on p. 80).

[W107] Martin Holland. *Cloud-Dienstleister OVH: Feuer Zerstört Rechenzentrum, Ein Weiteres Beschädigt*. 2021-03-10. URL: https://www.heise.de/news/OVH-Feuer-zerstoert-Rechenzentrum-in-Strassburg-ein-weiteres-beschaedigt-5076320.html (cit. on p. 87).

[W159] Joshua Lund. *Technology Preview for Secure Value Recovery*. 2019-12-19. URL: https://signal.org/blog/secure-value-recovery/ (visited on 2021-07-12) (cit. on p. 87).

[W171] MG Chemicals. *MG Chemicals Specialty Adhesives Catalog*. 2019. URL: https://www.mgchemicals.com/downloads/catalogs/Specialty%20Adhesives%20Catalogue%20Web.pdf (visited on 2021-07-08) (cit. on p. 86).

[W187] Lily Hay Newman. *Apple's T2 Security Chip Has an Unfixable Flaw*. 2020-10-06. URL: https://www.wired.com/story/apple-t2-chip-unfixable-flaw-jailbreak-mac/ (cit. on p. 75).

[W189] Karsten Nohl, Fabian Bräunlein, and dexter. *Shopshifting: The Potential for Payment System Abuse*. 2015-12-27. URL: https://media.ccc.de/v/32c3-7368-shopshifting#t=2452 (cit. on pp. 92, 95).

[W193] Gemalto NV. *SafeNet PCI-e HSM 6.2 Product Documentation: High Availability (HA) Overview*. 2015-12-18. URL: https://thalesdocs.com/gphsm/luna/6.2/docs/pci/Content/administration/ha/ha_overview.htm (visited on 2021-07-12) (cit. on p. 87).

[W194] Johannes Obermaier. *Physical Unclonable Functions: The Future Technology for Physical Security Enclosures?* 2019-08-24. DOI: 10.5446/43265 (cit. on pp. 86, 88, 91, 92).

[W205] Trevor Perrin. *The Noise Protocol Framework*. Version Revision 34. 2018-07-11. URL: http://noiseprotocol.org/noise.html (visited on 2021-07-13) (cit. on p. 100).

[W221] Root Zone KSK Operator Policy Management Authority. *Root Zone KSK Operator Key Management Procedure*. Version Version 3.4. 2021-09-22. URL: https://www.iana.org/dnssec/procedures/ksk-operator/KSK_Key_Management_Procedure_v3.4.pdf (visited on 2021-10-07) (cit. on p. 91).

[W247] Daniel Terdiman. *Aboard America's Doomsday Command and Control Plane*. 2013-07-23, 2013-07. URL: https://www.cnet.com/news/aboard-americas-doomsday-command-and-control-plane (cit. on p. 80).

Patent References

[P214] Mujib Rahman. "Optical Fiber Cable with Tampering Detecting Means". U.S. pat. Patent US4859024A. 1988-03-10 (cit. on p. 79).

References

[8] Nils Albartus et al. "DANA Universal Dataflow Analysis for Gate-Level Netlist Reverse Engineering". In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* 2020.4 (2020), pp. 309–336. DOI: 10.13154/tches.v2020.i4.309-336 (cit. on p. 75).

[14] Ross Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 1st ed. Wiley, 2020-12-22. DOI: 10.1002/9781119644682 (cit. on pp. 2, 22, 24, 36–38, 43, 64, 75, 77, 78, 92, 118).

[22] C. F. Beards. *Structural Vibration: Analysis and Damping*. Wiley, 1996. ISBN: 0-340-64580-6 (cit. on p. 89).

[24] Jean-Baptiste Bédrune and Gabriel Campana. "Everybody Be Cool, This Is a Robbery!" In: *Symposium Sur La Sécurité Des Technologies de l'information et Des Communications 2019*. 2019 (cit. on p. 93).

[30] David G. Boak. *A History of U.S. Communications Security, Volumes I and II*. Lecture Notes. US National Security Agency (NSA), 1973 (cit. on pp. 77, 78).

[36] Bertrand Campagnie. *Choose the Right Accelerometer for Predictive Maintenance*. Analog Devices, 2019 (cit. on p. 84).

[52] Jeremy Crawford. *Dungeons & Dragons - Player's Handbook*. Renton: Wizards of the Coast LLC, 2024. 384 pp. ISBN: 978-0-7869-6951-7 (cit. on p. 73).

[56] John C. Dixon. *The Shock Absorber Handbook*. Wiley, 2007. ISBN: 978-0-470-51020-9 (cit. on p. 89).

[58] Saar Drimer, Steven J Murdoch, and Ross Anderson. "Thinking inside the Box: System-Level Failures of Tamper Proofing". In: *2008 IEEE Symposium on Security and Privacy (Sp 2008)*. IEEE, 2008, pp. 281–295 (cit. on pp. 75, 78, 79).

[63] Maged Elsaid Elnady. "On-Shaft Vibration Measurement Using a MEMS Accelerometer for Faults Diagnosis in Rotating Machines". PhD thesis. University of Manchester, 2013 (cit. on p. 84).

[71] H. S. Fowler. *An Investigation of the Flow Processes in a Centrifugal Compressor Impeller*. National Research Council Canada, 1966. DOI: 10.4224/40003753 (cit. on p. 93).

[73] Jessie Frazelle. "Securing the Boot Process: The Hardware Root of Trust". In: *ACM queue : tomorrow's computing today* (2019-12-01). DOI: 10.1145/3380774.3382016 (cit. on p. 75).

[76] Yoshimitsu Fukushima and Teiji Tanaka. "A New Attenuation Relation for Peak Horizontal Acceleration of Strong Earthquake Ground Motion in Japan". In: *Bulletin of the Seismological Society of America* 80.4 (1990), pp. 757–783. ISSN: 0037-1106 (cit. on p. 89).

[88] A. German et al. "Event Data Recorders in the Analysis of Frontal Impacts". In: *Annual Proceedings of the Association for the Advancement of Automotive Medicine*. 51. 2007, pp. 225–243 (cit. on p. 89).

[94] Jan Sebastian Götte and Björn Scheuermann. "Can't Touch This: Inertial HSMs Thwart Advanced Physical Attacks". In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2022), pp. 69–93. DOI: 10.46586/tches.v2022.i1.69-93 (cit. on pp. 75, 194, 200, 208).

[118] Vincent Immler et al. "Secure Physical Enclosures from Covers with Tamper-Resistance". In: *IACR transactions on cryptographic hardware and embedded systems.* (2019). DOI: 10.13154/tches.v2019.i1.51-96 (cit. on pp. 3, 77–79, 82).

[122] International Atomic Energy Agency. *Safeguards, Techniques and Equipment*. Vol. 1. International Nuclear Verification Series. 2011. ISBN: 978-92-0-118910-3 (cit. on pp. 35, 36, 78).

[123] Phil Isaacs et al. *Tamper Proof, Tamper Evident Encryption Technology*. Surface Mount Technology Association / Surface Mount Technology Association, 2013 (cit. on pp. 77, 78, 85).

[128] Scott Johnson et al. "Titan: Enabling a Transparent Silicon Root of Trust for Cloud". In: *Hot Chips: A Symposium on High Performance Chips*. 2018 (cit. on p. 75).

[131] S. Graham Kelly. *Fundamentals of Mechanical Vibrations*. 2nd ed. McGraw-hill Series in Mechanical Engineering. McGraw-Hill, 1993. ISBN: 0-07-230092-2 (cit. on p. 89).

[141] Ivar Koene, Raine Viitala, and Petri Kuosmanen. "Internet of Things Based Monitoring of Large Rotor Vibration with a Microelectromechanical Systems Accelerometer". In: *IEEE access : practical innovations, open solutions* (2019). DOI: 10.1109/ACCESS.2019.2927793 (cit. on p. 84).

[142] Tony Kordyban. *Hot Air Rises and Heat Sinks: Everything You Know about Cooling Electronics Is Wrong*. ASME, 1998. ISBN: 978-0-7918-0074-4 (cit. on pp. 7, 86).

[143] Heinz Kreft and Wael Adi. "Cocoon-PUF, a Novel Mechatronic Secure Element Technology". In: *2012 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)* (2012). DOI: 10.1109/ahs.2012.6268655 (cit. on pp. 77, 79).

[195] Johannes Obermaier and Vincent Immler. "The Past, Present, and Future of Physical Security Enclosures: From Battery-Backed Monitoring to PUF-based Inherent Security and Beyond". In: *Journal of Hardware and Systems Security* 2 (2018), pp. 289–296. DOI: 10.1007/s41635-018-0045-2 (cit. on pp. 2, 37, 63, 78, 218).

[199] Diego Ongaro and John Ousterhout. "In Search of an Understandable Consensus Algorithm". In: *2014 USENIX Annual Technical Conference (USENIX ATC 14)*. Philadelphia, PA: USENIX Association, 2014-06, pp. 305–319. ISBN: 978-1-931971-10-2 (cit. on p. 87).

[223] Maruthi G. S. and Vishwanath Hegde. "Application of MEMS Accelerometer for Detection and Diagnosis of Multiple Faults in the Roller Element Bearings of Three Phase Induction Motor". In: *IEEE Sensors Journal* 16.1 (2016). DOI: 10.1109/JSEN.2015.2476561 (cit. on p. 84).

[233] Younes Shabany. *Heat Transfer: Thermal Management of Electronics*. CRC Press, 2009. ISBN: 978-1-4398-1468-0 (cit. on p. 86).

[236] Sean W Smith and Steve Weingart. "Building a High-Performance, Programmable Secure Coprocessor". In: *Computer Networks* 31.8 (1999-04), pp. 831–860. DOI: 10.1016/S1389-1286(98)00019-X (cit. on pp. 37, 64, 78, 95).

[251] Johannes Tobisch, Christian Zenger, and Christof Paar. "Electromagnetic Enclosure PUF for Tamper Proofing Commodity Hardware and Other Applications". In: *TRUDEVICE 2020: 9th Workshop on Trustworthy Manufacturing and Utilization of Secure Devices* (2020-03-13) (cit. on pp. 77, 79, 82).

[254] Timothy Trippel et al. "WALNUT: Waging Doubt on the Integrity of MEMS Accelerometers with Acoustic Injection Attacks". In: *2017 IEEE European Symposium on Security and Privacy*. IEEE, 2017, pp. 3–18 (cit. on p. 96).

[258] Hannes Tschofenig, Manuel Pegourie-Gonnard, and Hugo Vincent. "Performance of State-of-the-Art Cryptography on ARM-based Microprocessors". In: *NIST Lightweight Cryptography Workshop 2015*. 2015-07-21 (cit. on p. 100).

[266] Serge Vrijaldenhoven. "Acoustical Physical Uncloneable Functions". MA thesis. Technische Universiteit Eindhoven, 2004-10-01 (cit. on pp. 77, 79).

Chapter 5

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

We are as gods and might as well get good at it.

– *Stewart Brand [275]*

Contents

1 Introduction

This chapter is adapted from a paper written by me that will be presented by me at CHES 2026 [95].

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 [W126]. Other applications include card payment terminals where PCI PTS HSM standards [203] 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 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.

To do
Integrate new scope plots!

Inertial HSMs are one approach that enables the use of less expensive, commodity materials in high-security applications. Several other academic approaches exist that target low-cost [265, 263, 60, 264] or high-performance mesh monitoring [116, 117, 80]. Some academic works even try to replace the security mesh with entirely different tamper sensing primitives [241, 262]. 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

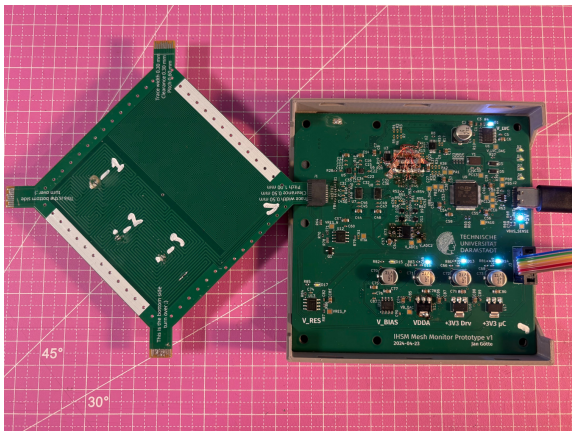

Figure 5.1: Measurement setup. Shown are the test specimen board on the left, and the frontend board with one of the four pulse amplifiers in the center. The frontend board is powered through a USB-C connection, and data is sent to a computer through a Single-Wire Debug (SWD) interface. The grid in the background has 10 mm pitch.

the mesh with a macro-scale Physical Unclonable Function (PUF) has also been researched [116, 241, 262], 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 chapter, 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 design directly applies to IHSM implementations, and complements the security offered by the IHSM's mechanical motion.

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 5.1.

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 [265, 263]. 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 [265, 263].

- 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 [13, 245]. 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 [14].

Patent literature gives a partial view of commercial developments in this area. Even in recent patents such as [P34, P191, P216, P272, P149, P45] 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 [116, 60, 264], 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 [241, 262]. 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. [116], Obermaier et al. [197], and Garb [80] 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 Physical Unclonable Function (PUF), combining tamper sensing with cryptographic key storage. In their design, the mesh consists of a crosshatch 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 [197] using an approach similar to Sato, Poupyrev, and Harrison [228], 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 [78].

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, [W]126, 203]. 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. [60] 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 [264] 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 [201]. 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. [265] and Vasile and Svasta [263], 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. A 2007 patent [P167] proposes the same delay-based approach.

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 [130]. 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 [130] trace back the style of four-diode balanced bridge sampling gate that we use to a vacuum tube implementation presented in Chance et al. [43]. 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 [111].

While initially equivalent time sampling was used to circumvent technological limitations, more recently it has also been used to achieve cost-optimized designs [110]. Going along similar principles, Polášek [208] 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

[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.

improves upon the double sampling design first presented by Houtman [110] 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. [25] 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 [185] show an equivalent time sampling TDR that uses specialized adjustable delay line ICs for pulse generation. Lee, Sung, and Park [147] 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. [253] show another FPGA-based TDR. Their system also uses a part from the same early FPGA family as Bencivenni et al. [25], 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 [110]. 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 $\frac{\text{V}}{\text{ns}}$ 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 [147], 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 [75, 178]. 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 [180] configuration measuring reflected signal components as well as using two or more ports measuring transmitted signal components [287]. Both Time Domain Reflectometry [75] and conventional frequency-domain VNA measurements [178] 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 [75] or as part of an FPGA gateware [178, 179]. 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 [178].

Variations of impedance analysis techniques have been demonstrated that detect changes inside individual chips using board-level measurements [158], that detect manipulatoins using non-contact near-field Radio Frequency (RF) measurements [225], that detect the mechanical preparation of a target chip for backside attacks using onboard measurements [179], and that adapt the technique as an offensive tool for side-channel analysis (SCA) attacks [176].

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 comple-

ment 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 chapter, 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 [116, 197, 80]. 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 chapter 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

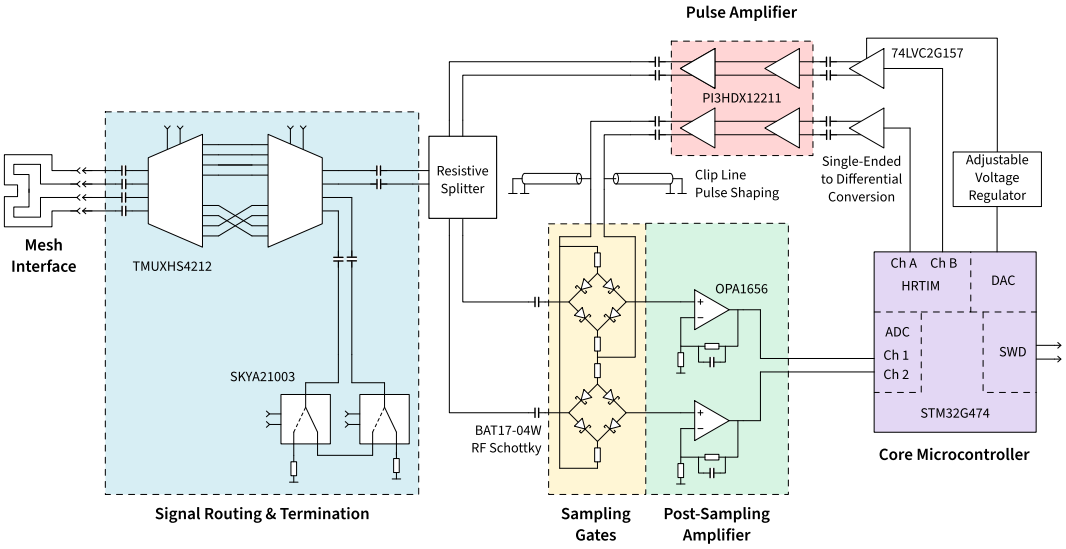

Figure 5.2: Block diagram of our prototype sampling TDR security mesh monitoring circuit.

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 5.2 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 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 [208, 110], 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

[2]Full schematics are available in the supplementary material of this thesis.

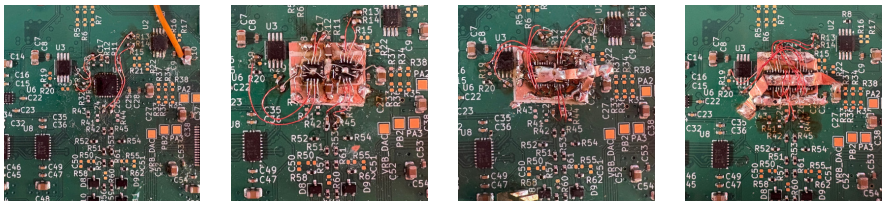

(a) 74LVC2G157 (b) MAX3748 (c) TDP0604 (d) PI3HDX12211

Figure 5.3: 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.

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 $t_r = t_f = 1.7\,\text{ns}$[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 [246] 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 chapter: A 74LVC2G157 standard logic IC, two HDMI/DisplayPort redrivers, PI3HDX12211 and TDP0604, as well as MAX3748, a limiting amplifier for optical networking. Figure 5.3 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

[3]Datasheet specification, when driving a 10 pF load [240].

times than oumicrocontroller [219].

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.

4.2 Cost Breakdown

Table 5.1 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] 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

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**

Table 5.1: Cost breakdown of our prototype design. Prices are listed at order quantity 1000 to make prices more comparable between distributors.

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 5.2, 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[31].

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.

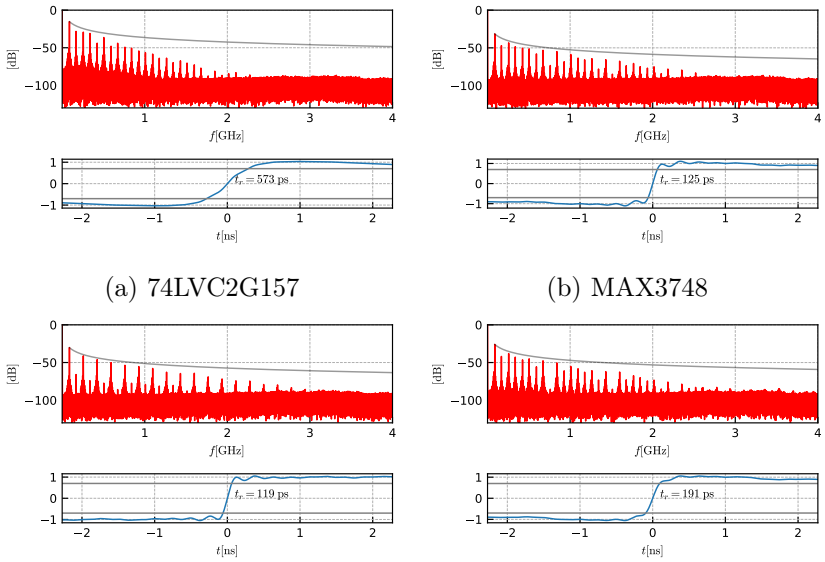

(a) 74LVC2G157 (b) MAX3748

(c) TDP0604 (d) PI3HDX12211

Figure 5.4: 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.

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), and through self-characterization of the circuit (Section 5.1). 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.

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

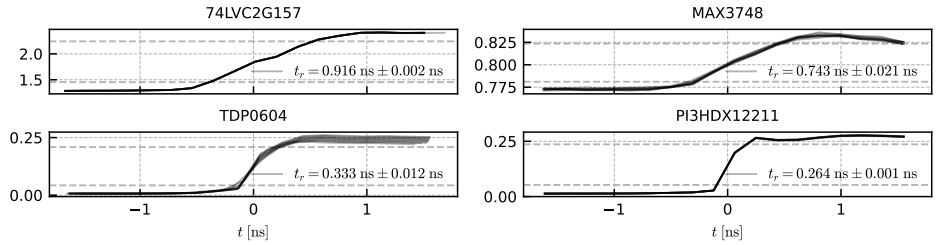

Figure 5.5: One edge of the stimulus pulse with no mesh connected measured by the board itself, using different amplifier ICs. For each IC, ten traces are shown. The vertical scale is in Volts at the sampling amplifier output.

26.5 GHz signal analyzer[5]. All measurements were taken with the prototype's mesh interface connected to the spectrum analyzer through a bias tee configured for DC blocking followed by a 20 dB attenuator for protection.

Figure 5.4 and Table 5.2 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.

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 5.5 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–

[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.

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 $\frac{\text{V}}{\text{ns}}$	1.89 $\frac{\text{V}}{\text{ns}}$	2.13 $\frac{\text{V}}{\text{ns}}$	2.25 $\frac{\text{V}}{\text{ns}}$

Table 5.2: 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 5.4, 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.

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 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 5.2 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 5.4 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.

Figure 5.6 shows the sampling and stimulus pulse edges measured using

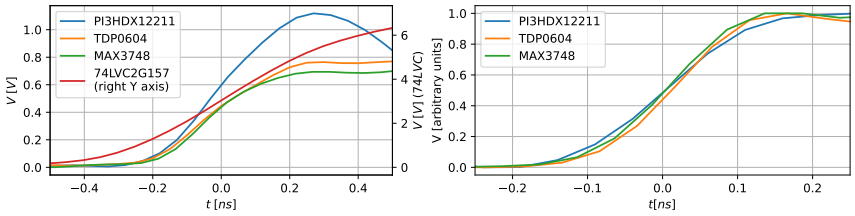

(a) Sampling pulse (b) Stimulus pulse

Figure 5.6: Oscilloscope measurements of the sampling pulse probed differentially (left) and of the stimulus pulse probed single-ended and normalized (right). The 74LVC pulse is plotted on the right Y axis in the left plot due to its large amplitude. In the right plot, it is not shown since our measurement setup did not allow for a measurement of this amplitude.

a Siglent SDS7404A 4 GHz oscilloscope. The stimulus pulse was directly measured single-ended, and the sampaling pulse was measured differentially through a Siglent SAP2500D 2.5 GHz active differential probe. These measurements support the conclusion from Figure 5.4 that in raw edge risetime, MAX3748 and TDP0604 perform fastest, with PI3HDX12211 being slightly slower. They also exhibit the large differences in amplitude that we expect cause the differences in actual measurement performance as shwon in Figure 5.5. Note that due to the differences in measurement methodology, a direct comparison of the rise times is not possible between these plots. The spectrum measurements do not convey amplitude information and discard low-frequency content, but due to the very large bandwidth of the spectrum analyzer used, they will represent the true risetime the closest. In both the self-characterization and the oscilloscope measurements, the displayed risetime is contaminated by the measurement system. In case of the self-characterization, the stimmulus rise time is folded into the measurement result, leading in the displayed risetime being slower by a factor of $\sqrt{2}$. Similarly, in the oscilloscope measurements, the combined risetime of the oscilloscope frontend and active probe contaminate the results.

We observed the best overall performance 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.

Mesh	1	2	3	4
Size	$35 \times 70\,\text{mm}$	$35 \times 70\,\text{mm}$	$35 \times 70\,\text{mm}$	$35 \times 70\,\text{mm}$
Area	$24.5\,\text{cm}^2$	$24.5\,\text{cm}^2$	$24.5\,\text{cm}^2$	$24.5\,\text{cm}^2$
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

Table 5.3: Specifications of mesh test specimens used in the experiments in this chapter. Approximate signal delays were calculated using wave velocity $v = \frac{c}{\sqrt{\epsilon_r}} \approx \frac{c}{2}$ [274] assuming $\epsilon_r \approx 4$ [184] for the test specimens' FR-4 substrate.

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 5.3 shows the design specifications. Each specimen contains four separate meshes on 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 5.7 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 5.7 by calculating speed of light in our mesh specimen's substrate based on them. The resulting measurements are shown in Table 5.4. All amplifier configurations yield comparable measurements of approximately $1.6\,\frac{\text{m}}{\text{s}}$, which corresponds with the expected signal propagation velocity in FR-4 PCB material of $1.5 \times 10^8\,\frac{\text{m}}{\text{s}}$ [274, 184].

The graphs in Figure 5.7 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 mon-

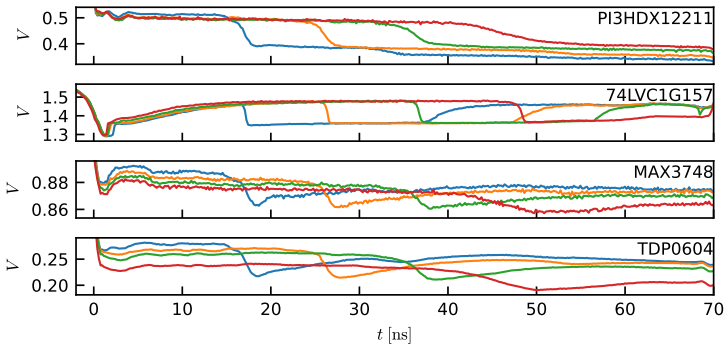

Figure 5.7: TDR responses captured by the microcontroller's internal ADCs with each of four candidate pulse amplifier ICs and four test meshes. The shown time range covers the primary reflection of the stimulus pulse's falling edge. For clarity, only one channel of the differential response is shown.

	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	$1.59 \times 10^8 \frac{\mathrm{m}}{\mathrm{s}}$
74LVC2G157	17.1 ns	26.4 ns	36.6 ns	48.2 ns	$1.55 \times 10^8 \frac{\mathrm{m}}{\mathrm{s}}$
MAX3748	17.2 ns	26.4 ns	36.6 ns	45.6 ns	$1.59 \times 10^8 \frac{\mathrm{m}}{\mathrm{s}}$
TDP0604	17.0 ns	26.2 ns	36.5 ns	45.8 ns	$1.59 \times 10^8 \frac{\mathrm{m}}{\mathrm{s}}$

Table 5.4: Speed of light and time offset calculated from delays read from the graphs in Figure 5.7. c is the speed of light determined by linear fit.

itoring, meshes should be broken up into segments that are multiplexed through signal switching.

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.

Interpreting these performance plots

Figure 5.8 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 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.

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

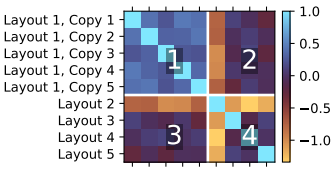

(a) Five copies of the same layout compared to four other layouts. FNR 18% at 0.1% FPR, CER=8.3%.

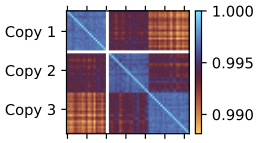

(b) Three identical copies, 20 measurements each. FNR 1.7% at 0.1% FPR, CER=1.1%.

Figure 5.8: Similarity matrices of measurement series on intact meshes.

Figure 5.8a 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 5.8a 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 5.8b 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 chapter, the presence of this effect indicates good performance of our design, and increases the detection efficiency of our approach.

Basic attacks

Figure 5.9 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 Error Rate.

Trace shortening

Figure 5.10 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-

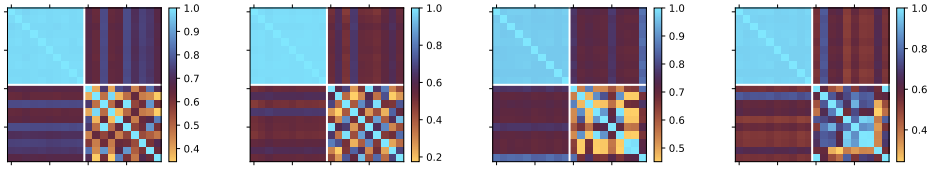

(a) One trace interrupted, p=0.3 mm. FNR 0.0% at 0.1% FPR, CER=0.0%.

(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 5.9: 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.

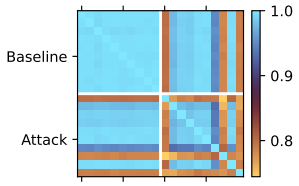

Figure 5.10: Similarity matrix of several mesh specimens that have one trace shorted to an adjacent location on the same trace. Classification FNR 23% at 0.1% FPR, CER=22%.

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.

Advanced attacks

Figure 5.11 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 5.11a), a soldering iron touching one mesh trace (Figure 5.11b), and a mesh where one trace has a $l = 30\,\text{mm}, d = 120\,\mu\text{m}$ piece of copper wire soldered to one trace (Figure 5.11a). 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.

[7]Part number Rigol PVP3150.

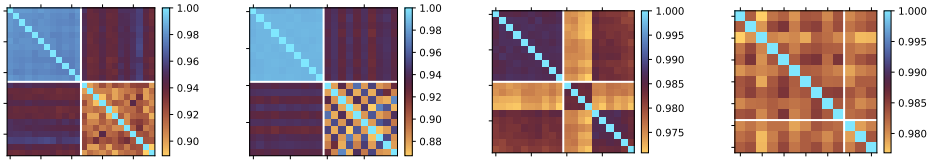

(a) Oscilloscope probe contacting mesh. FNR 0.0% at 0.1% FPR, CER=0.0%.

(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 5.11: Classifier performance under advanced attack scenarios.

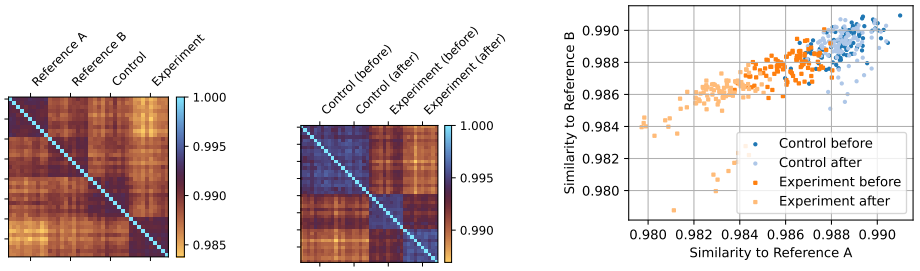

(a) Test boards before experiment.

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

(c) Trajectory of relative difference to reference specimens.

Figure 5.12: Classifier performance under a patching attack that bridges a short gap within a mesh trace using wire.

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 [117]. We performed such an attack on a set of 300 µm pitch meshes. Figure 5.13 shows our modification and the resulting change in the time-domain response.

Figure 5.12 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 5.12a shows the four samples before the attack, ex-

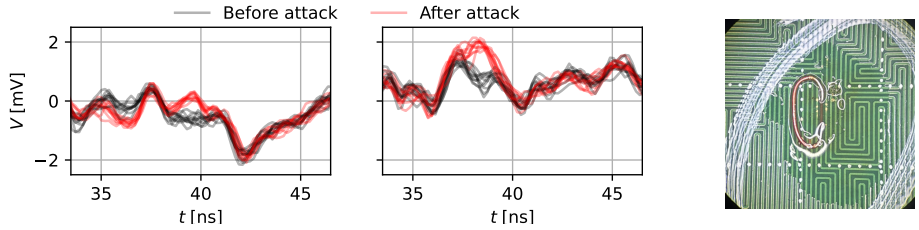

Figure 5.13: The mesh response under a manipulation attack patching across a drill location for a 300 µm drill, as captured by the microcontroller's ADCs. The mesh pitch is 300 µm. B-spline smoothing was applied for readability.

hibiting 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 5.12b compares both control and experiment samples before and after the attack, and shows a clear change in the experiment sample during the attack. Figure 5.12c 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 5.14 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 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, soldering for instance is highly detectable, while here we only benchmark a momentary snapshot after the patch was completed.

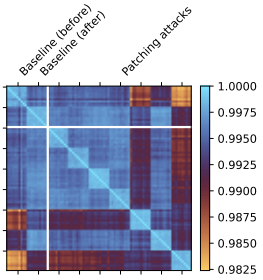

(a) Micro-soldering patching attack. FNR 71.5% at 0.1% FPR, CER=29%.

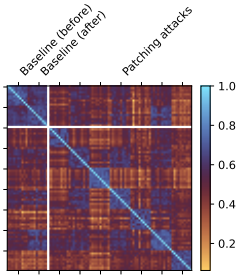

(b) *maximum* classifier variant. FNR 51.1% at 0.1% FPR, CER=15%.

Figure 5.14: Classification performance in a larger-scale experiment using 10 measurements each of 7 samples with traces patched through micro-soldering.

Environmental susceptibility

Figure 5.15 shows the results of a series of experiments evaluating the effect of environmental factors such as handling or electromagnetic interference on our measurements. Figure 5.15a shows our measurements exhibit little time drift (CER=60%). Figure 5.15b 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 5.15c, heating the mesh distors its measurements (FNR=0.6%, CER=0%). Figure 5.16 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 [226, 106]. 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 5.16 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,

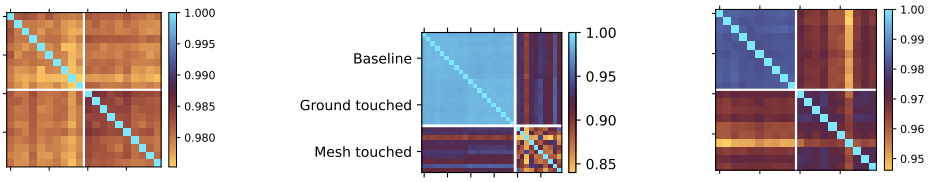

(a) Time drift (2.5h). FNR 100% at 0.1% FPR, CER=61%.

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

(c) Mesh heated (70 °C).

Figure 5.15: Classification results of the same mesh under various environmental factors.

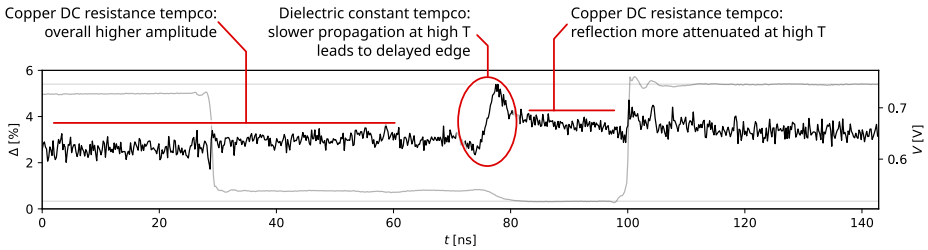

Figure 5.16: The effect of heating on a time-domain trace. One of 12 channels shown. Gray: Raw data. Black: Relative difference between hot and cool cases.

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 5.17 shows the result of these measurement series. For comparison, we included several measurements from Figure 5.14. 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%.

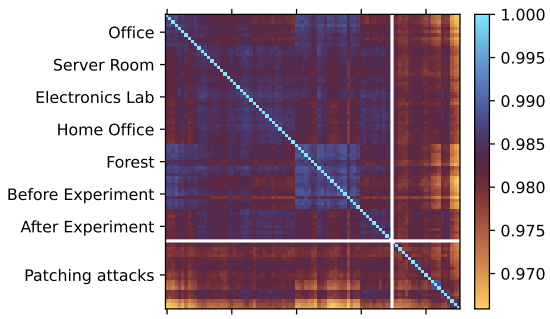

Figure 5.17: Classifier similarity scores of measurements in different environments, 10 measurements each. For scale, measurements from Figure 5.14 are included on the bottom/right. FNR 69.0% at 0.1% FPR, CER=22%.

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 [203] contains a useful framework for thinking about 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 benign 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 chapter could be used for other monitoring tasks from tamper sensing to system health monitoring. For instance, Vai et al. [262] 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 chapter, 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. The improved monitoring approach we presented in this chapter directly complements the IHSM concept we introduced in Chapter 4. Both designs can be combined into a joint system that provides a level of tamper resistance beyond the state of the art in both acadmic designs and in commercial offerings.

Web sources

[W126] *ISO/IEC 24759:2025*. ISO. URL: `https://www.iso.org/standard/82424.html` (visited on 2025-04-08) (cit. on pp. 5, 24, 115, 119).

Patent References

[P34] William L. Brodsky et al. "Tamper-Respondent Assembly With Flexible Tamper-Detect Sensor(s) Overlying In-Situ-Formed Tamper-Detect Sensor". U.S. pat. 10,327,329 B2. Internationnal Business Machines Corporation. 2019-06-18 (cit. on p. 118).

[P45] Raymond O. Chock and Mark Hess. "Point of Sale Terminal Having Pulsed Current Tamper Control Sensing". U.S. pat. 7551098B1. Zilog Inc. 2009-06-23 (cit. on p. 118).

[P149] Alan Henry Leek and Jace Hunter Hall. "Tamper Detection". U.S. pat. 10,925,154 B2. Texas Instruments Incorporated. 2021-02-16 (cit. on p. 118).

[P167] Noriaki Matsuno. "Protection Circuit for Semiconductor Device and Semiconductor Device Including the Same". U.S. pat. 7345497B2. Matsushita Electric Industrial Co Ltd. 2008-03-18 (cit. on p. 120).

[P191] John Norton. "Tamper Detecting Cases". U.S. pat. 10489614B2. Hewlett Packard Enterprise Development LP. 2019-11-26 (cit. on p. 118).

[P216] Mani Razaghi and Jesse Hill. "Tamper Detection System". U.S. pat. 10595400B1. Square Inc. 2020-03-17 (cit. on p. 118).

[P272] Erling Wesselhoff. "Tamper Responsive Sensor". U.S. pat. 10678957B2. Cryptera AS. 2020-06-09 (cit. on p. 118).

References

[1] (US) National Institute of Standards and Technology. *Security Requirements for Cryptographic Modules*. Federal Information Processing Standard (FIPS) 140-2. U.S. Department of Commerce, 2002-12-03. DOI: 10.6028/NIST.FIPS.140-2 (cit. on pp. 2, 4, 24, 115, 119).

[13] R. Anderson et al. "Cryptographic Processors-A Survey". In: *Proceedings of the IEEE* 94.2 (2006-02), pp. 357–369. DOI: 10.1109/JPROC.2005.862423 (cit. on p. 118).

[14] Ross Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 1st ed. Wiley, 2020-12-22. DOI: 10.1002/9781119644682 (cit. on pp. 2, 22, 24, 36–38, 43, 64, 75, 77, 78, 92, 118).

[25] 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 (2013-01-11), pp. 185–191. DOI: 10.1016/j.nima.2012.10.023 (cit. on p. 122).

[31] Eric Bogatin. *Signal and Power Integrity, Simplified*. Third edition. Boston: Prentice Hall, 2018. 958 pp. ISBN: 978-0-13-451341-6 (cit. on p. 131).

[43] Britton Chance et al., eds. *Waveforms*. Vol. 19. MIT Radiation Laboratory. New York: McGraw-Hill, 1949 (cit. on p. 121).

[60] François Dupont et al. "A Miniaturized and Ultra-Low-Power Tamper Detection Sensor for Portable Applications". In: *IEEE Sensors Journal* 22.5 (2022-03), pp. 4524–4533. DOI: 10.1109/JSEN.2022.3143656 (cit. on pp. 115, 118, 119).

[75] 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 (2018-10), pp. 1320–1324. DOI: 10.1109/TCSII.2018.2858798 (cit. on pp. 122, 123).

[78] Kathrin Garb et al. “FORTRESS: FORtified Tamper-Resistant Envelope with Embedded Security Sensor”. In: *2021 18th International Conference on Privacy, Security and Trust (PST)*. 2021 18th International Conference on Privacy, Security and Trust (PST). 2021-12, pp. 1–12. DOI: 10.1109/PST52912.2021.9647783 (cit. on pp. 118, 198).

[80] Kathrin A Garb. “Tamper-Sensitive Design of PUF-Based Security Enclosures” (cit. on pp. 3, 115, 118, 124, 198).

[95] Jan Sebastian Götte and Björn Scheuermann. “High Fidelity Security Mesh Monitoring Using Low-Cost, Embedded Time Domain Reflectometry”. In: *Transactions on Cryptographic Hardware and Embedded Systems*. Conference on Cryptographic Hardware and Embedded Systems 2026. Vol. 2026/1. IACR, 2026 (cit. on p. 115).

[106] Scott Hinaga et al. “Thermal Effects on PCB Laminate Material Dielectric Constant and Dissipation Factor”. In: IPC Apex Expo. 2010 (cit. on p. 143).

[110] Hubert Houtman. “1-GHz Sampling Oscilloscope Front End Is Easily Modified”. In: *Electronic Design* 48.19 (2000-09-18), pp. 175–176. ISSN: 0013-4872 (cit. on pp. 121, 122, 127).

[111] *HP 187B Dual-Trace Vertical Amplifier Operating and Service Manual*. Hewlett-Packard Company, 1962 (cit. on p. 121).

[116] 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)*. 2018 IEEE International Symposium on Hardware Oriented Security and Trust (HOST). 2018-04, pp. 49–56. DOI: 10.1109/HST.2018.8383890 (cit. on pp. 3, 37, 115, 116, 118, 124).

[117] Vincent Immler et al. “Secure Physical Enclosures from Covers with Tamper-Resistance”. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2018-11-09), pp. 51–96. DOI: 10.46586/tches.v2019.i1.51-96 (cit. on pp. 37, 115, 141, 198).

[130] M. Kahrs. "50 Years of RF and Microwave Sampling". In: *IEEE Transactions on Microwave Theory and Techniques* 51.6 (2003-06), pp. 1787–1805. DOI: 10.1109/TMTT.2002.806934 (cit. on p. 121).

[147] 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)*. 2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515). Vol. 2. 2003-10, 1219–1223 Vol.2. DOI: 10.1109/NSSMIC.2003.1351912 (cit. on p. 122).

[158] 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. DOI: 10.1007/978-3-030-88428-4_34 (cit. on p. 123).

[176] 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, 2023-11-21, pp. 1675–1689. DOI: 10.1145/3576915.3623092 (cit. on p. 123).

[178] 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* (2022-11-29), pp. 301–325. DOI: 10.46586/tches.v2023.i1.301-325 (cit. on pp. 122, 123).

[179] Tahoura Mosavirik and Shahin Tajik. "BackMon: IC Backside Tamper Detection Using On-Chip Impedance Monitoring". In: *Proceedings of the 2024 Workshop on Attacks and Solutions in Hardware Security*. CCS '24: ACM SIGSAC Conference on Computer and Communications Security. Salt Lake City UT USA: ACM, 2024-11-19, pp. 68–77. DOI: 10.1145/3689939.3695784 (cit. on p. 123).

[180] Tahoura Mosavirik et al. "Silicon Echoes: Non-Invasive Trojan and Tamper Detection Using Frequency-Selective Impedance Analysis". In: *IACR Transactions on Cryptographic Hardware and Embed-*

ded Systems 2023.4 (4 2023-08-31), pp. 238–261. DOI: 10.46586/tches.v2023.i4.238-261 (cit. on p. 123).

[184] 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 (1989-03), pp. 287–292. DOI: 10.1007/BF02657420 (cit. on pp. 136, 163).

[185] Catalin Negrea and Marius Rangu. "Sequential Sampling Time Domain Reflectometer". In: 2009 15th International Symposium for Design and Technology of Electronics Packages (SIITME). 2009-09, pp. 367–371. DOI: 10.1109/SIITME.2009.5407341 (cit. on p. 122).

[197] 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, 2018-06-24, pp. 1–6. DOI: 10.1145/3195970.3195976 (cit. on pp. 118, 124, 198).

[201] 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 (1977-02), pp. 616–620. DOI: 10.1109/TNS.1977.4328751 (cit. on p. 120).

[203] PCI Security Standards Council. *Payment Card Industry PIN Transaction Security Hardware Security Module Modular Security Requirements*. Version 4.0. 2021-12 (cit. on pp. 115, 119, 145).

[208] Petr Polášek. "Reflektometr v Časové Oblasti". MA thesis. 2020-01-30 (cit. on pp. 121, 127).

[219] Renesas Electronics Corporation. *Application Note AN-224: ALVC/LVC Logic Characteristics and Applications*. 2019 (cit. on p. 129).

[225] Maryam Saadat Safa and Shahin Tajik. "Near-Field Microwave Sensing for Chip-Level Tamper Detection". In: *Sensors* 25.13 (2025-07-05), p. 4188. DOI: 10.3390/s25134188 (cit. on p. 123).

[226] Pankaj Sagar and Kashif Akber. "Studies on Temperature Dependent Dielectric Properties of Some Insulators down to Liquid Helium Temperatures". In: *Cryogenics* 141 (2024-07-01), p. 103865. DOI: 10.1016/j.cryogenics.2024.103865 (cit. on p. 143).

[228] 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, 2012-05-05, pp. 483–492. DOI: 10.1145/2207676.2207743 (cit. on p. 118).

[240] ST Microelectronics. *STM32G474xB/C/E Datasheet*. 2021-11 (cit. on p. 128).

[241] Paul Staat et al. "Anti-Tamper Radio: System-Level Tamper Detection for Computing Systems". In: *2022 IEEE Symposium on Security and Privacy (SP)*. 2022 IEEE Symposium on Security and Privacy (SP). 2022-05, pp. 1722–1736. DOI: 10.1109/SP46214.2022.9833631 (cit. on pp. 115, 116, 118).

[245] Mark Tehranipoor et al. *Hardware Security Primitives*. Cham: Springer International Publishing, 2023. DOI: 10.1007/978-3-031-19185-5 (cit. on p. 118).

[246] Tektronix Inc. *Tektronix S-6 Sampling Head Instruction Manual*. 1982-09 (cit. on p. 128).

[253] Dennis Trebbels et al. "Miniaturized FPGA-Based High-Resolution Time-Domain Reflectometer". In: *IEEE Transactions on Instrumentation and Measurement* 62.7 (2013-07), pp. 2101–2113. DOI: 10.1109/TIM.2013.2245190 (cit. on p. 122).

[262] Michael Vai et al. "Secure Architecture for Embedded Systems". In: *2015 IEEE High Performance Extreme Computing Conference (HPEC)*. 2015 IEEE High Performance Extreme Computing Conference (HPEC). 2015-09, pp. 1–5. DOI: 10.1109/HPEC.2015.7322461 (cit. on pp. 115, 116, 118, 146).

[263] 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)*. 2017 IEEE 23rd International Symposium for Design and Technology in Electronic Packaging (SIITME). 2017-10, pp. 175–178. DOI: 10.1109/SIITME.2017.8259885 (cit. on pp. 115, 117, 120).

[264] 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). 2019-10, pp. 212–215. DOI: 10.1109/SIITME47687.2019.8990877 (cit. on pp. 37, 115, 118, 119).

[265] 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)*. 2017 40th International Spring Seminar on Electronics Technology (ISSE). 2017-05, pp. 1–6. DOI: 10.1109/ISSE.2017.8000987 (cit. on pp. 115, 117, 120).

[274] H.A. Wheeler. "Transmission-Line Properties of Parallel Strips Separated by a Dielectric Sheet". In: *IEEE Transactions on Microwave Theory and Techniques* 13.2 (1965-03), pp. 172–185. DOI: 10.1109/TMTT.1965.1125962 (cit. on p. 136).

[275] *Whole Earth Catalog Spring 1969*. Point Foundation, 1969. 132 pp. (cit. on p. 113).

[287] 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. DOI: 10.1109/TIFS.2023.3285490 (cit. on p. 123).

Chapter 6

Rotation-Invariant Envelope Power Supply

> A knot is never "nearly right"; it is either exactly right or it is hopelessly wrong, one or the other; there is nothing in between. This is not the impossibly high standard of the idealist, it is a mere fact for the realist to face.
>
> – *Clifford Ashley [16]*

Contents

A central engineering challenge in inertial HSMs is transferring power and data between the payload and the rotating mesh cage (cf. Chapter 4). Industrially, power and data transfer through rotating joints is usually done using slip ring assemblies. A slip ring consists of one or more contacts that wipe on a rotating circular surface. Industrially, metal spring contacts plated with hard gold or other common surface coatings are used for transferring small currents and data signals, and carbon brushes are used for higher currents. Slip rings are widely used in motors and other rotating machinery.

For use in IHSMs, slip rings have several limitations. First, they are complex precision-machined components and thus are rather expensive. Beyond cost, they also have performance limitations. Generally, slip rings are most well-suited to slow rotation, as high rotation increases the wear of the contacts. The design target of 1000 rpm we use in IHSMs are at the upper end of what commercial slip rings usually support. A third disadvantage is that they are sensitive, and any misalignment or contamination by dust can increase wear and cause intermittant contact.

An IHSM's data link can easily be realized using optical communication. Although power transfer using light is also possible—and we have in fact demonstrated it in our first prototype IHSM—it comes at the disadvantage of a heavy rotating assembly since large solar cells are needed, and it has poor end-to-end efficiency. For the large-scale meshes needed in a high-performance IHSM tailored to SMPC applications, we engineered a better solution: A rotation-invariant inductive Wireless Power Transfer link.

While Wireless Power Transfer (WPT) is widely used and can be implemented in many different ways [18, 21, 53, 67, 148, 151, 162, 177, 181, 182, 285]. Most WPT variants link the primary and secondary side primarily through the magnetic component of the electromagnetic field, and coils are used as the transmitting and receiving antenna. Such *inductive* WPT uses low frequency, which reduces circuit complexity, and it is well-suited for transferring high power across short distances. The electronic realization of a WPT link is usually similar to that of a DC/DC converter, except that in place of the inductor or flyback transformer, the pair of transceiver coils is used. Compared to a flyback transformer, the WPT link's transceiver coil pair has a lower coupling coefficient that varies with distance.

A challenge in WPT links is the strong dependency between link inductor coupling coefficient and distance. In a naïve implementation that uses

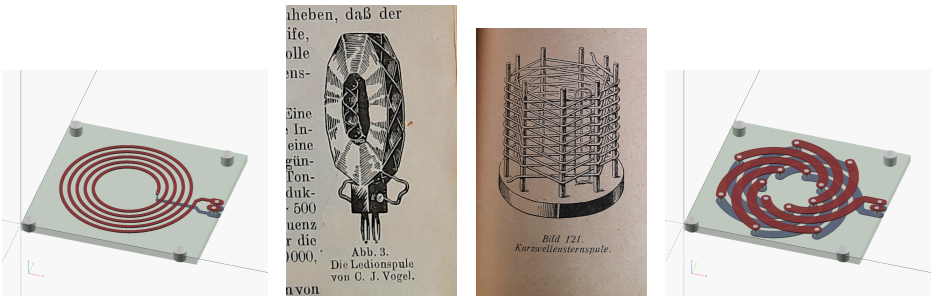

(a) A classic planar spiral inductor

(b) A honeycomb coil in Saacke [224]

(c) A basket-woven coil in Klein [137]

(d) Our proposed inductor layout

Figure 6.1: Illustration of our proposed inductor layout compared to contemporary conventional planar inductors and honeycomb as well as basket-woven coils from the early days of wireless radio.

the link coils as a simple transformer, link efficiency would drop sharply with distance. To decrease the impact of this distance dependency, almost all WPT implementations combine the transceiver coils with capacitors to form a pair of tuned tank circuits that are driven like they would be in a resonant converter. Like in resonant converters, a variety of topologies such as series, parallel, or series-parallel LC are used for these tuning circuits.

To do
Not final graphics. Get proper scans for camera-ready version

1 Construction Approach

In the WPT link powering the rotating mesh of an IHSM presents an unusual set of constraints, which does not seem to be addressed adequately in the existing literature on inductive WPT yet. To reduce the need for custom-wound inductors, we settled on using a planar inductor implemented in a Printed Circuit Board (PCB). Such planar PCB inductors are limited by the structure size limits of the PCB process, resulting in rotational asymmetry due to the trace width. Planar inductors are usually considered approximately axisymmetric. In our application, we found that the field asymmetry in feasible PCB inductors is large enough that axial rotation of two such inductors results in an oscillation of their coupling coefficient that leads to voltage ripple on the secondary side, especially when the coils are misaligned.

The large centrifugal acceleration on an IHSM mesh prohibits the use of batteries or liquid electrolyte capacitors on the rotating part, and makes

heavy components such as large Multilayer Ceramic Capacitors (MLCCs) and ferrite-core inductors challenging to balance. As a result, the secondary-side voltage ripple poses a significant issue since the conventional ways of efficiently filtering such ripple through large bypass capacitors or through a secondary-side switchmode power supply are difficult to implement due to their mass.

In other inductive WPT systems, this issue is mitigated by one of several factors: First, for this effect to matter in the first place, the two coils have to be rotating with respect to one another. In ferrite core inductors, the core is the major factor shaping the magnetic field and evens out the small effect of winding asymmetry. In wire-wound inductors, the often higher turn count and the tightly packed, circular wires render this effect negligible. Finally, the output ripple caused by this oscillation can be filtered through a voltage regulator or by using a large decoupling capacitor on the secondary side if the application can accomodate such components on the rotating part.

While there exist a corpus of prior work focusing on efficient power transfer between two coils whose position relative to one another cannot be precisely controlled as is the case in wireless phone charging systems as well as in proposed WPT electric vehicle chargers [151, 181], it is generally assumed that the two coils remain quasi-stationary with respect to one another.

There exists a body of work on inductive power transfer through rotating joints but here the focus often lies on higher power budgets than our application requires, which in practice requires more space and a ferrite or laminated iron core [67, 237, 267]. Often, these rotating joint WPT systems use coaxial structures, but segmented approaches exist, too [268, 279, 278, 152]. In lower-power applications, segmented approaches are more common. A key challenge in segmented approaches is the reduction of secondary-side ripple induced when the segments' alignment changes throught one revolution [284], which usually requires additional secondary-side circuitry. In this work, we introduce a planar coil topology for WPT through a rotating joint using a single planar PCB coil on both the transmitting and the receiving side that improves rotation ripple at low turn counts.

1.1 Twisted inductors

To solve these issues, we propose a layout for circular PCB inductors that uses a number of series-connected interleaved spirals to achieve a topological

equivalent to a torus knot from mathematical knot theory. Our layout twists the inductor's windings around one another by connecting the interleaved spiral segments with a ring of vias each on the inside and outside of the inductor's windings. Our approach provides better performance beyond our particular use case, and improves over conventional contemporary planar inductors applying similar principles to those which inspired the polygonal basket-woven air coils used in early radio sets. We show that we can layout a twisted inductor for any number of layer inversions that is co-prime to the inductor's turn count. Our approach opens up a design space for inductor layouts that interpolate between planar spiral inductors on one end, and planar toroidal inductors on the other end. Our approach thus generalizes a super-set to a number of previous approaches to the design of planar inductors.

We observe that in high-frequency applications, a moderate number of layer inversions increases the spacing between the beginning and end of the inductor's conductor, where the majority of the inductor's AC current flows. This decreases the parasitic capacitance of the inductor and increases its Self-Resonant Frequency (SRF), raising its maximum possible operating frequency and improving its efficiency at lower operating frequencies.

1.2 Contributions

Our contributions on this matter include:

- We introduce twisted inductors, a planar inductor layout that improves rotational symmetry in WPT through rotating joins, and promises improved high-frequency behavior in other applications.
- We provide detailed instructions for the construction of such layouts, including a mathematical analysis of the available parameter space.
- We provide an analytical model of inductance and DC equivalent series resistance of our scheme.
- Validating our scheme, we provide laboratory measurements of the basic parameters of 39 test specimens comparing our scheme to conventional layouts.
- We further present the results of Finite Element Method (FEM) simulations to validate our inductance and ESR approximations.

- Finally, to analyze the degree of rotational symmetry in our proposed scheme, we provide the results of a large number of automated measurements of coupling between pairs of inductors under various rotations, offsets, distances and load conditions.

2 Related Work

2.1 Inductive WPT in Practice

Inductive WPT has been proposed in a large number of scenarios [285, 182], each of which comes with a set of unique constraints. When WPT is used to charge an electric toothbrush, the implementation cost of the system is critical, while efficiency and total power output are of little concern. Mechanically, in an electric toothbrush's charging system, the position and spacing of the transmitter and receiver coils can easily be controlled down to millimeter precision.

In contrast to this, wireless smartphone charging is a much more demanding application. Here, the total cost of the system is only secondary, but the receiver's form factor is critical, and total power output as well as efficiency become major objectives. At the same time, in wireless smartphone charging, position tolerances are very coarse, and the two coils in the charging base and in the phone may be positioned more than a centimeter off-axis, with a gap of several millimeters and potentially not even in parallel planes.

Power transfer across large distances is even more of a concern in implantable medical devices [177]. Where a wireless phone charger must be able to bridge distances of a few millimeters, an implantable medical device might be situated underneath several centimeter of tissue and bones. At the same time, cost is of (almost) no concern in this medical application, which enables the use of complex manufacturing techniques, customized electronic components and exotic materials.

While all of the aforementioned applications transfer somewhere between a few hundred milliwatts and several watts of power, at the other end of the spectrum there is a large body of research suggesting the use of inductive wireless power transfer for the charging of electric vehicles (EVs) [151, 181]. In this application, the wireless power transfer system usually replaces the conventional wired charging connector, which improves the systems' user experience given the strong force required to seat or unseat

these rather large connectors, as well as the heft of the required water-cooled cables. In this application, size is of little concern, but at charging rates up to tens of kilowatt, efficiency becomes critical.

2.2 Core materials in WPT

Across application areas, air-core inductors are often used for WPT since in most applications, an air gap of several millimeters or more is expected [53]. Especially in low-power application such as mobile device charging, the size and weight of ferrites is an obstacle to their use, and at lower power levels losses are less of a concern.

A common way to use ferrites in WPT applications is by magnetically shielding the inductor's back side with a ferrite plate such that the field does not extend beyond the coil's back side, thereby increasing the intended mutual inductance while simultaneously reducing eddy current losses when the WPT coils are placed near metal objects [21, 148, 183]. Similar to how the trace layouts of planar WPT coils are optimized to improve power transfer efficiency, the layout of ferrite components has been proposed for optimization [21].

2.3 PCB inductor design for wireless power transfer

Today, air-core inductors are the standard solution in inductive WPT links. Since in most WPT applications an air gap of several millimeters between the sending and receiving assemblies is expected, adding a ferrite core does not result in a large improvement in coupling. Instead, the impact of this misalignment is reduced by maximizing the area of the air-core inductors used, or by tiling multiple inductors [53, 268, 283].

WPT inductors tend to be mostly planar coils with only a few layers, so implementing them in a PCB process seems natural. Using a PCB for the inductor has the potential to reduce implementation cost since PCBs are cheap, and they can also serve as structural support. However, implementing inductors in PCBs has several disadvantages. First, due to the limited layer count of common PCB processes and due to structure size limitations, the number of windings that can be fit into a given volume is much lower than in wire-wound inductors. Second, due to a PCB's copper layers being thin compared to its dielectric substrate[1] PCB inductors tend to have

[1] common values are 15 µm to 30 µm copper thickness and 600 µm to 1600 µm substrate thickness

poor DC resistance, albeit the thin copper layer decreases skin effect losses compared to a solid, round conductors of the same cross-sectional area. However, PCBs can still not approach the performance of litz wire used in high-frequency WPT coils, which commonly use wire diameters in the range of tens of micrometer [286]. Lope et al. [155] and Nomoto et al. [190] propose a mitigation that aims to emulate a litz wire's structure in large, high-current PCB inductors, but their mitigation is heavily limited by the structure size achievable in common PCB manufacturing processes [188].

A further factor that limits the high-frequency performance of PCB inductors is distributed capacitance. Not only does a large air coil exhibit more parasitic capacitance than an equivalent, smaller ferrite-core inductor simply due to its size, when implemented in a PCB process a large fraction of the electrical fields responsible for this capacitance pass through the PCB's substrate, not air. The relative permittivity ϵ_r of common PCB substrates typically lies in the range of 4 to 5 [184], which increases the distributed capacitance compared to a pure air-core inductor by approximately that same factor.

2.4 Planar Inductors in RFIC Design

Beyond WPT, planar inductors are commonly used in radio frequency integrated circuits (RFICs). In RFIC design, the major challenges are area optimization and precisely predicting the inductor's characteristics during the design phase. Common optimizations include applying a variable trace pitch [156] and variable trace width [112].

In RFICs, inductors are commonly designed as *balanced* inductors with a grounded central node. Such designs interleave two counter-wound planar spiral inductors on the same layer with the help of some jumper connections on a second layer [55, 166]. The use of such designs in RFIC design is mainly focused on their electrical symmetry, so that the two input ports can be fed with a fully differential signal, with the inductor loading both driver outputs equally across the inductor's frequency range.

2.5 A Brief Historical Diversion on Basket-Woven Air Coils

Since the early days of radio engineering, the parasitic capacitance of inductors has been a point of concern [186, 70]. Going back to the early days of wireless telegraphy after the turn of the twentieth century, coils with high inductance were needed for the construction of both transmitters and

receivers, but the ferrites that would later permit their compact construction were still being developed. The ferromagnetic core material of choice back then was laminated iron, which was only useful at low frequencies due to eddy current losses. As a result, the inductors in radio circuits of the era were often constructed as air-core coils. While air-core inductors are immune to core saturation, the poor magnetic permeability of air necessitates a large number of wide turns of wire to reach useful inductance values, which for reasons of practicality or leakage inductance often could not be wound as a single layer cylindrical coil. This could be resolved by winding an inductor with many turns on multiple layers, which improves compactness and leakage inductance, but this in turn gives rise to increased distributed capacitance as now turns with a large voltage differential are layered right on top of each other.

Before the invention of ferrites, a number of ways were devised to decrease distributed capacitance in multilayer inductors. These methods can be divided into two general categories: Optimizing the connecting order of turns and optimizing the winding schema of turns. Both aim at increasing spacing between parts of the coil that have a large voltage differential.

The connecting order of turns was optimized at the assembly level by stacking coils in a particular way [70] and at the component level by winding coils in a particular way to minimize the voltage differential between adjacent turns—a technique that is still used to this day [154]. The main winding optimization in the first category concerns winding the turns of a cylindrical multilayer inductor not layer by layer, but instead layering them diagonally, effectively connecting adjacent turns in a diagonal zigzag pattern. Then as now, wound inductors applying this technique were not feasible to manufacture reliably by machine, but the technique can be closely replicated in PCB inductors as shown in Lee, Su, and Ron Hui [146]. The main limiting factors in a PCB implementation are the requirement for a large number of vias inside the inductor's turns limiting the achievable turn count[2] and increasing equivalent series resistance (ESR) through the thin trace sections that are necessary to accomodate the via structure, as well as the layer pairing limitations when blind vias are used in multilayer PCBs.

This lack of a way to wind high frequency inductors with a machine

[2]In PCBs, as opposed to integrated circuits (ICs), vias limit the achievable turn count when they need to be placed in-line inside the turns as opposed to on the inside or outside because a PCB's minimum trace/space widths are usually much smaller than the smallest feasible via, consisting of a minimum-size drill surrounded by a minimum-size annular ring.

led to the creation of a number of related winding schemes that include honeycomb and basket woven coils [64, 68, 137, 169, 192, 242, 276, 288]. The simplest such winding technique is the universal winding as described in depth by Querfurth [211]. In a simple, cylindrical wire-wound inductor, the windings are laid down one right next to the other, until the end of the winding area is met, where the winding direction is reversed. One layer of such windings forms a helix whose pitch is equal to the wire diameter. A universal winding uses the same helical scheme reversing at the coil ends, but uses a helical pitch larger than the wire diameter to form a structure similar to a spool of sewing thread.

Other winding techniques include honeycomb and basket woven coils, some historic examples of which are shown in Figure 6.1. In a honeycomb coil, like in an universal winding, subsequent winding layers are wound at a criss-cross pattern. The characteristic feature of honeycomb coils is that the winding machine is adjusted to produce large air gaps between adjacent windings, resulting in a three-dimensional rhomboid pattern that is vaguely reminiscent of a honeycomb's structure.

In basket-woven coils, a mandrel consisting of an odd number of sticks pointing either radially or axially is used, and the wire is woven between adjacent sticks in an alternating direction. While visually similar to honeycomb coils, this winding technique is more suited to homebrew construction and less amenable to mass production by machine. In axially basket-woven coils, the mandrel can be pulled out after the coil is finished. Like honeycomb coils, the resulting structure can be made mechanically stable with some lacquer, with the turns carrying the layers where they cross.

Both construction techniques apply similar principles to those leading to the improved high-frequency behavior of twisted inductors that we describe in this chapter.[3]

3 Twisted Inductor Design

In this section, we present a detailed derivation of the layout of twisted inductors. We approach this layout by construction. Let us first consider a simple, planar, circular spiral coil with a fixed pitch. We will ignore

[3]Interestingly, the winding schemes of both honeycomb and basket-woven coils are also governed by the same coprimality condition between the number of turns and the number of inversions within each turn that we describe for our twisted inductors below, although we could not find an example in historic literature where this condition was explicitly stated [64, 137, 276, 211].

trace width for now, and consider the trace a thin wire. We will assume the inductor's ports are both located on the positive x-Axis on top of one another on different layers, which also helps to minimize the loop area of the inductor's connections.

The trace trajectory of a standard planar spiral inductor can be parameterized in polar coordinates r, φ based on an Archimedean spiral:

$$r = a \cdot \varphi \tag{6.1}$$

An Archimedean spiral defined this way always starts at the origin, and it continues to infinity. Let us re-parameterize this spiral to a curve parameter t with range $[0, 1]$, such that $t = 0$ corresponds to the start of the inductor and $t = 1$ corresponds to its end. As is customary in PCB inductors, we place the inductor's start on its outer circumference.

To improve layer utilization, a common technique in PCB inductor design is to use both layers of the PCB for the inductor's spiral trace, instead of only using the bottom layer for a straight jumper trace. Using both layers this way allows for wider traces, which lowers resistive losses. We can accomodate this optimization in our definition by re-defining our normalized radius to allow both positive and negative values, defining negative values to designate traces on the PCB's bottom layer as follows. Figure 6.2 shows both a simple and a two-layer spiral inductor in the first two columns.

Let n be the turn count of our inductor. The resulting parametrization is:

$$\begin{aligned} \varphi &= 2\pi n t \\ r &= r_1 + |1 - 2t| \, (r_2 - r_1) \end{aligned} \tag{6.2}$$

The resulting spiral trace starts at radius r_2 on the positive x axis, and spirals inward until it meets r_1, where the sign indicates a layer change, and the trace reverses to continue back to r_2 on another layer. In its PCB realization, at r_1, a via would be placed to connect the end of the spiral trace to a jumper trace on the other layer of the PCB leading back to the start.

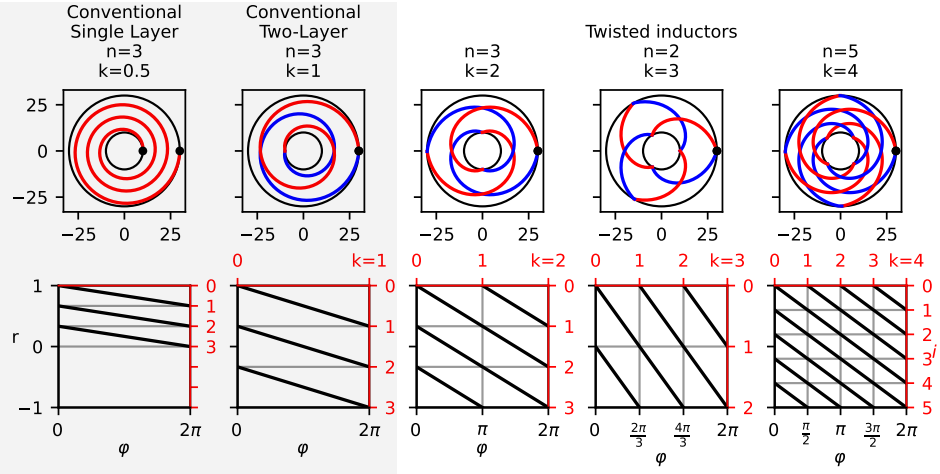

Figure 6.2: Inductor layouts for several sets of turn count n and inversion count k. The top row shows the actual trace layout in cartesian coordinates, the bottom row visualizes the winding schema.

3.1 From Spiral to Twisted Inductor

Extending the above parametrization of a spiral inductor's layout, we propose planar *twisted inductors* based on two core observations:

Observation 1.

When using an archimedean spiral, multiple such spirals using the same pitch can be interleaved by spreading out their start and end points at regular angular intervals.

Observation 2.

In a two-layer spiral inductor (Figure 6.2), we can adjust the turn count of the pair of traces to move the end point of the bottom layer trace anywhere on the inductor's outer radius.

Setting the inversion count to $k = 1$ in our proposed scheme yields the conventional two-layer counterwound scheme [154, 239, 146].

Combining these two observations, we find that by choosing a number k of inversions, i.e. layer jumps, that is coprime to the number of total turns of the inductor n, we achieve a layout where all k pairs of top and bottom-layer traces naturally connect in series, with the resulting spirals on the top and bottom layers interleaving cleanly. Figure 6.2 shows a layout with $n = 3$ turns with both a single inversion ($k = 1$), which results in a conventional two-layer inductor, and with $k = 2$ inversions, creating two interleaved spirals on both the top and the bottom layer of the PCB. Figure

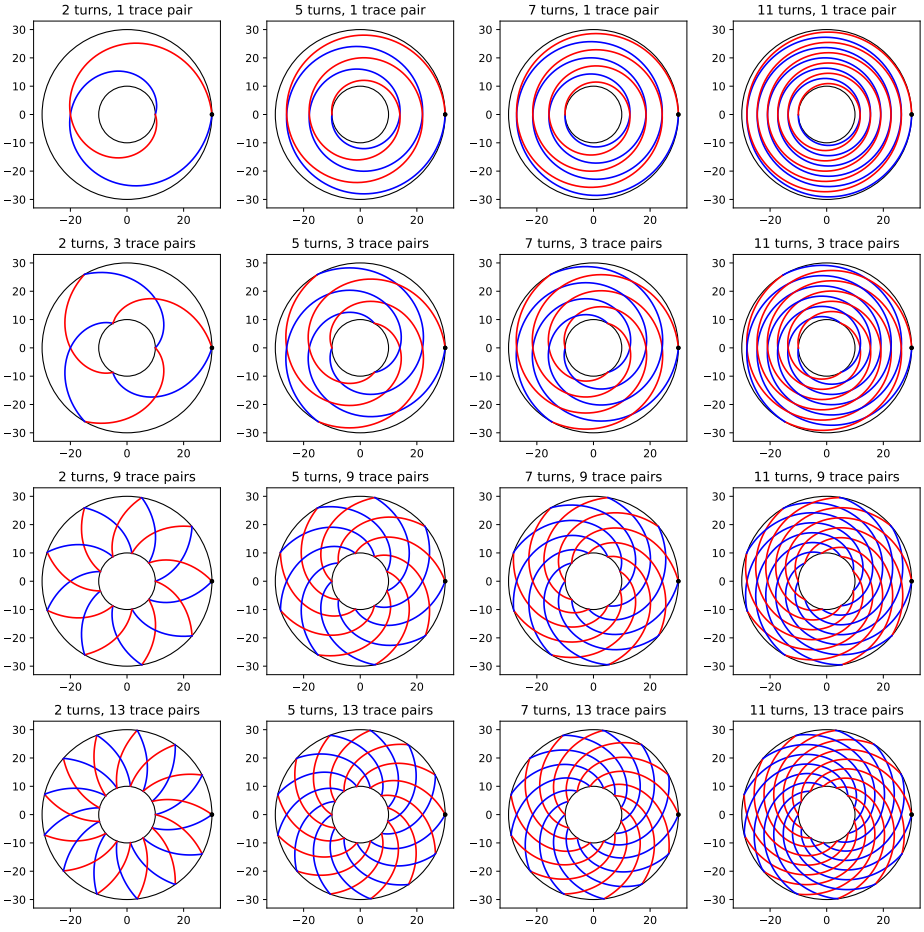

Figure 6.3: Layout examples for a number of combinations of turn count n and inversion count k. Note that in this illustration we chose values for n and k such that all pairs are coprime.

6.3 shows additional layout examples for other values of n and k. For $k = \frac{1}{2}$, we get a standard single-layer planar spiral inductor for any turn count n, and for $k = 1$ we get a standard two-layer planar spiral inductor for any turn count n. In this chapter, we will call all layouts with $k \geq 2$ *Twisted Inductors*. The coordinate description of Equation 6.3 thus becomes:

$$\begin{aligned} \varphi &= 2\pi n t \\ r &= r_1 + \left|1 - (2kt \mod 2)\right| (r_2 - r_1) \end{aligned} \tag{6.3}$$

Topologically, the shape of our inductors can be described as a (k, n)-torus knot. From knot theory, we know that such a torus knot exists if and only if both n and k are co-prime. Figure 6.2 illustrates a derivation of the coprimality requirement. If we plot the spiral in polar coordinates on a cartesian plot we observe that for a n-turn coil with k inversions, the trace crosses the φ axis once for each inversion, wrapping around r. Likewise, it crosses the r axis once for each turn of the inductor, wrapping around φ. Based on this, we can re-label the angular axis in steps from 0 to k, and re-label the radial axis in steps from 0 to n. Labelling the new angular axis i and the new radial axis j, in the resulting integer lattice, the trace has slope 1. We can state the trace's trajectory as a function of a curve parameter $t \in [0, nk]$ as $f(t) = (i, j) = (t \mod n, t \mod k)$. To produce a valid inductor, the trace must not intersect anywhere. Thus, the system of congruences

$$t \equiv i \mod n \tag{6.4}$$

$$t \equiv j \mod k \tag{6.5}$$

must have a unique solution $t \in [0, nk]$ for all (i, j). This statement corresponds exactly to the Chinese Remainder Theorem, which states that this solution is unique if and only if k and n are coprime.

In the following paragraphs, we will derive analytical expressions for Ohmic resistance and inductance of inductors derived under this schema.

Ohmic Resistance The arc length l of a spiral can be calculated from its turn count n and the average of its inner and outer diameter $\frac{2r_1+2r_2}{2} = r_1+r_2$ as $l = n\pi\,(r_1 + r_2)$. Since going from a standard inductor to a twisted

inductor does not change its turn count nor its dimensions, the combined arc length of all traces of the twisted inductor does not change. Twisted inductors require two additional vias per inversion, which will increase DC resistance slightly, but the contribution of these vias will remain small in practical applications since the overall number of vias is still no more than a couple per turn, and since each via only bridges the short distance between the inductor's layers.

As a general expression, for a standard or twisted inductor with turn count n and twist count k, given via resistance R_{via} we derive a first order approximation of the inductor's DC resistance as follows.

$$R_L = n\pi\frac{r_1 + r_2}{2} + (2k - 1)\, R_{\text{via}} \tag{6.6}$$

Inductance Even for geometrically simple inductors, analytically calculating their inductance is a surprisingly hard problem whose complexity quickly escalates when geometrically complex inductors are analyzed, when realistic wire shapes as opposed to thin wire or current sheet approximations are used, and when taking into account differing magnetic permeabilities of air or dielectrics and core materials. Instead of precise analytical models, a number of approximations are commonly used. A commonly referenced approximation for the inductance of planar spiral inductors is given by Mohan et al. [175], whose current-sheet approximation for circular planar spiral inductors we will use here to estimate our inductor's inductance. The current-sheet approximation from Mohan et al. [175] reads:

$$L = \frac{\mu n^2 d_{\text{avg}} c_1}{2}\left(\ln\left(c_2/\rho\right) + c_3\rho + c_4\rho^2\right) \tag{6.7}$$

In this equation, c_{1-4} denote four empirically determined coeficcients that are specific to the coil's shape. The values for circular coils are $c_{1-4} = (1.00, 2.46, 0.00, 0.20)$. μ is the magnetic permeability of air (for an air-core inductor), n is the number of turns, d_{avg} is the *average* turn radius, i.e. $d_{\text{avg}} = 2\frac{r_1+r_2}{2} = r_1 + r_2$. $\rho = \frac{r_2-r_1}{r_2+r_1}$ is the planar spiral inductor's *fill ratio*. The fill ratio encodes the fact that the inductor's turns have less flux linkage the closer to the inductor's center we get. While turns close to the outside have good flux linkage due to their inner area overlapping well with that of other turns, turns close to the center not only have a loop area that is only a fraction of that of turns further outwards, the closer we get to the center, the larger is also the fraction of the field lines returning as leakage

flux on the outside of the inner turn that pass through the inner part of turns further outwards, flipping the sign and contributing *negative* mutual inductance.

As Equation 6.7 approximates the inductor's whole set of windings as a single, uniform current sheet, the turn count only appears as a single factor of n^2 in the equation, with ρ and c_{1-4} correcting for the inductor's geometry. To account for twisted inductors, we can separate the inductor into a set of $2k$ simple planar spiral inductor *branches* that are connected in series by the twisted inductor's vias. Compared to a simple spiral inductor, for each branch, the inductance according to Equation 6.7 stays the same except that the factor n^2 drops to $\left(\frac{n}{2k}\right)^2$ because the n windings are evenly distributed across the $2k$ branches. Let us now make two assumptions. First, we will assume that the flux linkage between both sides of the inductor is approximately one. This assumption is grounded in the fact that for practical designs, the substrate thickness will be small compared to the inductor's diameter. Second, we will for now ignore the spiral inductor's field asymmetry and assume that the flux linkage between two intertwined branches on the same side of the substrate is approximately one. In our measurements below we show that for simple spiral inductors this asymmetry, while problematic in our application, is small in absolute terms, and grows smaller with increasing turn count.

Based on these two assumptions, we can model the twisted inductor as a set of $2k$ series-connected spiral inductors that are perfectly coupled, with full flux linkage. This results in the total series inductance gaining back the factor $\frac{1}{2k^2}$ that each branch lost, resulting in identical inductances for a simple planar spiral inductor and a twisted inductor with the same size and turn count according to Equation 6.7. This approximation introduces an error due to the imperfect flux linkage between the two sides of the substrate, and between two spiral branches located at an angular offset from each other. In our experiments, we found that for our test inductors, compared to inductances measured with an LCR meter, this error is below 10 % for $n = 5$ turns or more, and for our test samples matches the performance of Equation 6.7 for the simple planar spiral inductor case.

3.2 CAD Integration

To allow for easy design with twisted inductors and to speed up the laboratory prototyping we performed for this chapter, we created a tool that

generates arbitrary twisted inductor layouts, and that is able to output these layouts as PCB footprint files for the open source KiCad EDA CAD tool [W134]. We integrated the ESR and inductance approximations as derived above with our tool, so that it provides immediate design feedback when generating inductors. In order to minimize ESR and maximize PCB area utilization, we made the tool automatically calculate the largest possible trace width when given a minimum clearance specification.

To handle outputting PCB geometry in a format that can be read from KiCad, we utilized the open source EDA file format library *gerbonara* [87]. To support the FEM simulations that are described in the next section below, our tool contains functionality to map gerbonara's geometry representation into that of gmsh [89], the FEM mesher that we chose to interface with Elmer FEM [222].

4 FEM Simulation

To validate our analytical approximations, we performed a series of FEM simulations in Elmer FEM. For a number of inductor layouts, we performed simulations to determine ohmic resistance and inductance. Due to limitations in our gmsh/Elmer toolchain, we were unable to run simulations for parasitic capacitance and self-resonance, or for coupling behavior of coil pairs. We found that for these cases which require larger, more complex meshes, gmsh would frequently crash during meshing, and where we were able to produce meshes, Elmer would only converge for some of them. While these are problems that can be solved through either a more skillful description of the problem in gmsh and Elmer, or by using more robust software such as Simulia CST, we decided to instead experimentally measure these quantities instead (cf. Section 5). While our measurements only cover a small number of inductor samples, their results are more reliable than results from FEM and can serve as a baseline for future work on such simulations.

We conducted our FEM simulations as follows:

Ohmic Resistance In Elmer FEM, we can use the built-in joint static current and joule heating solver to determine the ohmic resistance at a given current.

Inductance We let Elmer determine inductance by first using its coil solver to determine the volumetric current density in our mesh given a test current, then applying its magnetodynamics solver to solve the electromagnetic field. Elmer provides routines to derive the total magnetic field energy U_{mag} from an EM field solution. Since we have only our inductor under test inside the simulation volume, with test current I_{test}, we can then derive the inductor's inductance according to the well-known relation [168]:

$$L = \frac{2 \cdot U_{\text{mag}}}{I_{\text{test}}^2} \tag{6.8}$$

5 Experimental Validation

To experimentally validate our design with real-world inductors, we produced 24 test coupons with a number of variations of twisted inductors with winding count n between 1 and 25, and twist count ranging from $k = \frac{1}{2}$ (simple single-sided spiral inductor) to $k = 37$. All test inductors had an inner diameter of 15 mm and an outer diameter of 35 mm corresponding to the space available in our IHSM implementation.

5.1 Inductance and DC resistance

We measured the inductance and DC resistance of each test coupon using a Keysight U1733C LCR meter at 100 kHz for inductance and a Keysight 34465A multimeter in four-wire configuration for DC resistance. We further determined the self-resonant frequency of each inductor using a LiteVNA64 handheld vector network analyzer. The results of our measurements are shown in Table 6.1.

We found our inductance approximation to be accurate within 10 % and our ESR approximation to be accurate within 20 % for inductors with three turns or more. For lower turn-count inductors, inductance measurements are difficult because the small absolute inductances involved are easily disturbed by stray inductances, and ESR measurements are affected by contact and trace resistance even when measurements are taken in four-wire mode.

In accordance with our design intuition, we found that for high turn count inductors, the doubled trace width that is afforded by splitting a simple spiral inductor across two PCB layers in any two-layer configuration improves ESR by approximately a factor of two. Going from a simple single-layer spiral inductor to a simple two-layer spiral inductor ($k = 1$), we

observe that the resulting inductance decreases by up to 15 %. We suspect that the main factor leading to this decrease is radial magnetic flux leakage through the PCB material between the inductor's layers. Comparing simple two-layer inductors with $k = 1$ to the twisted inductors with larger k values that we propose in this chapter, we observe almost identical performance for $k > 1$ with decreases of less than 0.5 % going from $k = 1$ to $k = 3$ irrespective of turn count. From these measurements we can conclude that the flux linkage of twisted inductors almost perfectly matches that of simple two-layer inductors.

Finally, we decided to evaluate the high-frequency performance of twisted inductors. It is well-known that self-resonant frequency decreases when going from a single-layer spiral inductor to a two-layer spiral inductor while keeping inductance and dimensions constant [282]. Our measurements show this effect, with it being more pronounced with higher turn count. Intuitively, this makes sense if we consider the mechanism of inductor self-resonance. The primary contributor to self resonance, particularly in higher turn count inductors, is capacitive coupling between the inductor's windings. In a single-layer spiral inductor, this effect gets partially mitigated since the strongest coupling exists between adjacent windings, which here have only a small voltage differential as only a fraction of the inductor's total voltage appears across each winding. Compared to this, when the inductor is constructed as a simple two-layer inductor with $k = 1$, now the start and end windings of the inductor, which have the highest voltage differential, are located right on top of each other with the substrate in between. Making things worse, common PCB substrates have a relative permittivity much larger than air (usually around 4).

We observe that this decrease in high-frequency performance is eventually counteracted by increasing inversion count k. While our test samples focused on smaller turn counts, we observe a notable increase from a self-resonant frequency of 8.9 MHz for a standard $n = 25, k = 1$ inductor to 10.6 MHz for $n = 25, k = 13$. Prompted by this observation, we produced another set of 15 samples focusing on this aspect. We report our results of this investigation in the following section.

In conclusion to the above measurement results, we observe that twisted inductors *improve* high-frequency performance compared to simple two-layer inductors while closely matching them in ESR and inductance. While they peform worse than simple single-layer inductors in high-frequency per-

formance, the increased trace width that two-layer inductors allow for lowers resistive losses by approximately a factor of four. In applications where resistive losses lead to the choice of a two-layer inductor, twisted inductors provide improved high-frequency performance at no additional cost and without compromising other performance parameters.

5.2 Inductance and Frequency Behavior of Larger Coils

To investigate the high-frequency behavior of twisted inductors further, we produced and measured 15 additional sample inductors that were larger (up to 90 mm outer diameter) and that had a higher turn count (up to 53) compared to our initial set of samples. The parameters of these new samples and our measurement results are shown in Table 6.2. In these results, we can identify three clear trends. First, the ESR of twisted inductors is generally poorer when compared to two-layer spiral inductors. This increase in ESR is due to the large number of vias used in these sample inductors. It should be noted that while twisted inductors have worse ESR compared to conventional two-layer inductors, in our first set of test coupons we saw that their ESR is still better than that of a single-layer inductor because the traces can be made wider. Our second observation is that in every set of samples from this second run of physically larger inductors, twisted inductors outperform conventional planar inductors in self-resonant frequency by a considerable margin with an increase in SRF of up to 58 % from our $d_2 = 65\,\text{mm}$ sample going from $k = 1$ to $k = 100$.

Our third observation is that unlike in the smaller inductors from Table 6.1, in these larger instances, twisted inductors show increased inductance by approximately 3.7 % for our smallest samples, and 6.5 % for our largest samples. This behavior indicates that large twisted inductors indeed behave like a combination between a conventional planar spiral inductor and a conventional planar toroidal inductor. Comparing the magnitude of this increase with the measurements listed in Table 6.2 for planar toroidal inductors, we see that this effect exceeds what one would reach by a simple series configuration of both styles of inductor, indicating a contribution from flux linkage.

5.3 Coupling and its Sensitivity to Radial Offset

To evaluate twisted inductors in our WPT application, we measured the variation of the coupling between a pair of inductors using an automated

Parameters		Design values				Simulation results				Measurements		
n	k	L [µH]	Error [%]	R [Ω]	Error [%]	L [µH]	Error [%]	R [Ω]	Error [%]	L [µH]	f_{res} [MHz]	R [Ω]
1	3	0.03	−93.1	0.0095	−49.9	0.039	−43.6	0.008	−78.8	0.056	**465.07**	**0.0143**
1	4	0.03	−103.4	0.0108	−38.6	0.040	−47.5	0.008	−87.5	**0.059**	460.08	0.015
1	5	0.03	−89.7	0.0123	−35.3	0.041	−34.1	0.009	−84.4	0.055	460.08	0.0166
2	1	0.12	−28.4	0.0253	−12.1	0.127	−17.3	0.024	−18.3	0.149	**245.51**	**0.0284**
2	3	0.12	−31.0	0.0270	−7.9	0.128	−18.8	0.025	−16.4	**0.152**	240.52	0.0291
2	5	0.12	−26.7	0.0299	−0.2	0.130	−13.1	0.027	−11.1	0.147	225.5	0.03
3	1	0.26	−10.0	0.0454	−1.6	0.262	−9.5	0.044	−4.8	**0.287**	**145.71**	0.0461
3	4	0.26	−9.6	0.0479	5.0	0.265	−7.9	0.046	1.1	**0.286**	**145.71**	**0.0455**
5	1	0.73	4.5	0.0755	−3.1	0.670	−3.4	0.074	−5.1	**0.693**	61.345	0.0778
5	3	0.73	4.3	0.0763	4.7	0.671	−3.4	0.074	1.8	**0.694**	**70.285**	0.0727
5	7	0.73	4.4	0.0802	16.2	0.675	−2.8	0.077	12.7	**0.694**	68.05	**0.0672**
10	1	2.90	6.3	0.2513	7.6	2.700	−0.7	0.250	7.1	**2.718**	24.076	0.2322
10	3	2.90	6.4	0.2520	10.5	2.700	−0.5	0.250	9.8	2.714	**28.571**	0.2255
10	7	2.90	6.4	0.2554	16.9	2.700	−0.5	0.252	15.8	2.713	28.072	**0.2122**
25	1	18.15	6.7	1.8843	9.7	16.900	−0.2	1.900	10.4	16.938	8.84	1.7024
25	3	18.15	6.8	1.8851	13.2	N/A	N/A	N/A	N/A	16.919	8.595	1.636
25	13	18.15	6.7	1.9016	18.9	16.900	−0.2	1.900	18.8	16.931	**10.555**	**1.5429**
25	37	18.15	6.0	2.0197	15.9	17.100	0.2	2.000	15.1	**17.066**	10.31	1.698

Table 6.1: Inductor sample design parameters and measured characteristics. All inductors have outer diameter 35 mm and inner diameter 15 mm. The missing values in the simulation results columns result from the solver failing to converge. Bolded values highlight the best performing coil of each turn count. Shaded rows indicate conventional two-layer planar inductors ($k = 1$).

d_1 [mm]	d_2 [mm]	n	k	L [µH]	R_{ESR} [Ω]	f_{Res} [MHz]	C_{p} [pF]
25	40	1	150	5.00	11.0	N/A	N/A
25	40	53	1	120	**19.6**	18.0	0.65
25	40	53	50	121	22.6	**27.5**	**0.28**
25	40	53	100	123	26.9	26.5	0.29
25	40	53	150	**125**	33.2	24.0	0.35
50	65	1	300	10.2	21.9	N/A	N/A
50	65	53	1	270	**35.7**	10.0	0.94
50	65	53	100	272	41.9	**15.8**	**0.37**
50	65	53	200	277	50.1	13.3	0.52
50	65	53	300	**280**	65.0	13.8	0.48
75	90	1	480	17.3	35.5	N/A	N/A
75	90	53	1	441	**50.7**	7.00	1.17
75	90	53	160	444	60.8	**10.0**	**0.57**
75	90	53	320	461	76.2	8.75	0.72
75	90	53	480	**470**	92.9	8.00	0.84

Table 6.2: Parameters and measurement results of a set of larger sample inductors. Bold values indicate best performance at a given size. Shaded rows indicate conventional planar toroidal ($n = 1$) or two-layer planar spiral inductors ($k = 1$).

measurement setup consisting of a 3D gantry built from an old 3D printer, with a fourth rotation axis provided by a small servo that allows us to position two inductor test coupons at arbitrary offsets and angles to one another.

To do
pics of 3d printer test setup

To approximate our application, we loaded the secondary inductor with a 10 Ω resistor while providing a signal at a 300 kHz carrier frequency to the primary inductor from a Siglent SDG6022X function generator as shown in Figure 6.4. We measured both the input and output voltages of the coupled inductor pair using Keysight 34465A multimeters in AC Root Mean Square (RMS) mode.

Figure 6.5 shows the ratio between input and output voltage of our test link for a set of three-turn inductors with multiple inversion numbers k when one inductor is rotated. In practical WPT setups, the transmitter and receiver coils are rarely aligned perfectly, so we show measurements across a range of radial offsets. In line with our inductance measurements, coupling is lower at $k > 0$ compared to a single-layer spiral inductor. Across one revolution, we find that the single-layer spiral inductor exhibits the most voltage ripple, with simple two-layer inductors with $k = 1$ already improving

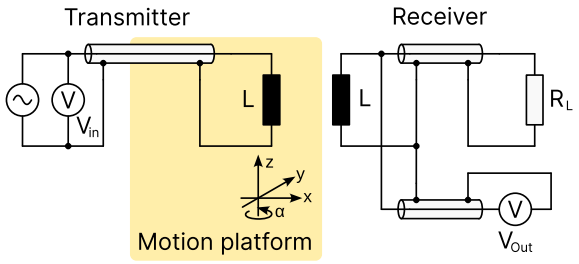

Figure 6.4: The test schematic used in all measurements. For direct coupling factor measurements, the load resistor was disconnected. We measure voltage at the output of the function generator to account for drop in its internal output resistance.

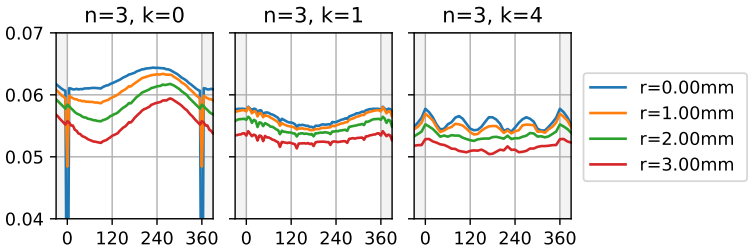

Figure 6.5: RMS output voltage of the test circuit from Figure 6.4 for three pairs of matching inductors with one inductor rotating w.r.t. the other. The inductors have $n = 3$ turns each and $k = \frac{1}{2}$, $k = 1$, and $k = 3$, respectively. For each k, voltage curves are plotted for a number of different radial offsets between the two inductor's centers.

ripple. For k above 1, ripple amplitude stay sconstant, but energy is shifted into higher frequencies that are easier to passively filter on the WPT link's secondary side in our application.

Expanding our measurements in the previous section, we performed a series of measurements rotating both inductors. In these measurements, the coils' distance is fixed 1 mm and the radial offset is set to a worst-case value of 4 mm. Figure 6.7 shows the normalized output voltage of a WPT link made from three-turn inductors with rotation of one inductor shown on the horizontal axis, and the rotation of the other shown on the vertical axis.

We performed similar measurements on 24 of our test coupons at 1 mm and 4 mm radial offsets. Figure 6.6 shows the combined results of these measurements, with worst-case voltage variation plotted across inversion count k for multiple turn counts n and radial offsets r. In this graph, we see that twisted inductors improve ripple compared to conventional designs, even at a low inversion count such as $k = 3$.

Concluding our measurements, we achieved our primary objective of reducing coupling variation under rotation, with twisted inductors ($k > 1$) improving over conventional two-layer spiral inductors, which perform better than simple single-layer spiral inductors. This improvement is greatest for inductors with low turn count and consequentially coarse pitch, as their turns deviate the furthest from a set of ideal, concentric circles.

6 Future Work

Our derivation of twisted inductors opens up a space for future research. On the practical side, as part of our inductor design tool, we extended the EDA file format library gerbonara with code to automatically map gerbonara's geometry description to the gmsh FEM mesher. This code may be of independent interest since it allows for the extraction of FEM meshes from not just individual planar components, but PCBs in any file format supported by gerbonara such as KiCad's native file format, as well as the Gerber file format supported by the majority of EDA tools.

On the theoretical side, the fact that our twisted inductor model generalizes both one- or two-layer planar spiral inductors as well as planar toroidal inductors would make the deduction of key parameters such as inductance and distributed capacitance by mathematical analysis or by finite element methods interesting. Furthermore, the precise contribution of vias to the

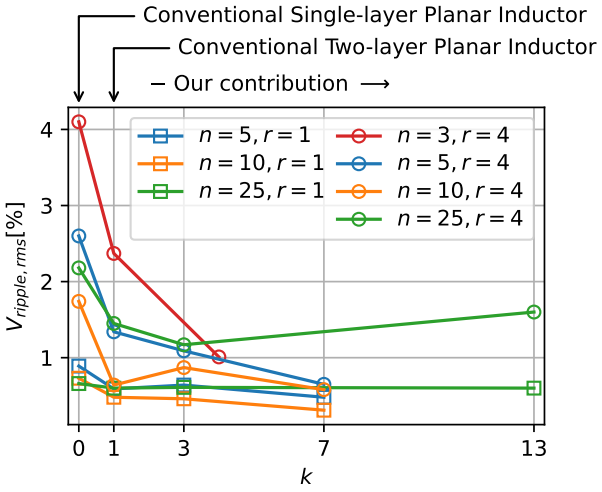

Figure 6.6: RMS Voltage ripple in a model rotating WPT setup with $R_L = 10\,\Omega$ as a percentage of total RMS output voltage, plotted against inductor inversion count k. Measurements were taken with a number of different coils with turn count n between a single turn and 25 turns. Measurements were taken at two different radial coil offsets of $r = 1\,\text{mm}$ and $4\,\text{mm}$. Coil distance was $d = 1\,\text{mm}$ in all cases. The shaded area indicates conventional coil layouts, with the remainder of the plot showing twisted inductors.

twisted inductor's parasitics is interesting, especially for layouts with large values of inversion count k. We suspect that via influence will be frequency dependant as vias and traces have distinct DC resistances, and skin effect will affect both to a differring extent.

7 Conclusion

In this chapter, we introduced a novel layout approach for planar, multi-layer inductors. Our *twisted* inductors generalize several types of conventional planar inductors including conventional single- or two-layer planar spiral inductors as well as planar toroidal inductors. For inversion count parameter $k \geq 2$, twisted inductors produce magnetic field distributions that have better rotational symmetry along the inductor's main axis compared to either conventional single- or two-layer planar spiral inductors, which yields lower output ripple in WPT through rotating joints and enables the use of smaller and lighter secondary-side circuitry, improving efficiency.

Furthermore, besides the advantages twisted inductors show in our particular application, we found that our sample twisted inductors have up

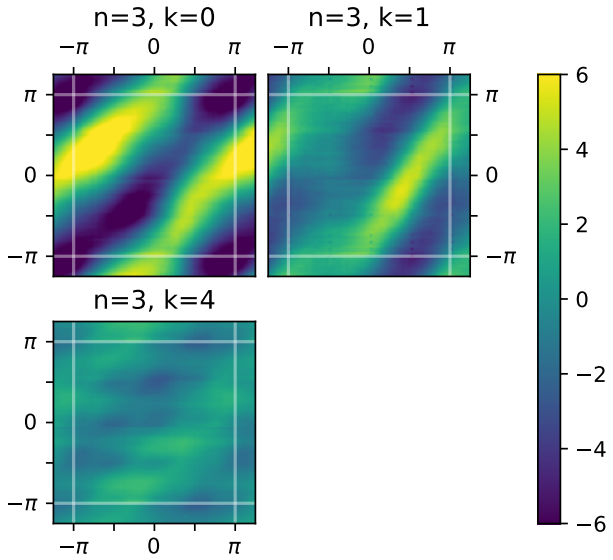

Figure 6.7: RMS ripple magnitude as a percentage of mean RMS output voltage, plotted against the rotation of each of the two inductors. The two coils were kept at a constant 4 mm radial offset, and the output coil was loaded with a 10 Ω load. All RMS ripple plots in this chapter share the same color scale to allow for visual comparison. This figure shows four variants of 3-turn coils, plots for $n = 5$ can be found in Figure 6.10 and plots for $n = \{10, 25\}$ in Figures 6.8 and 6.9.

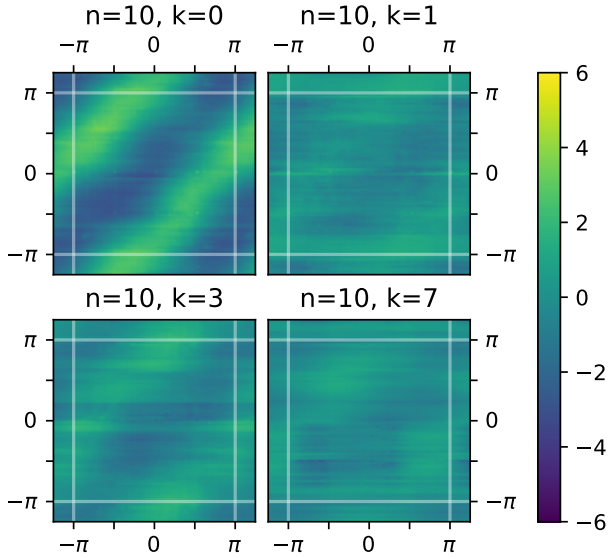

Figure 6.8: RMS ripple magnitude as shown in Figure 6.7 for four different 10-turn coils.

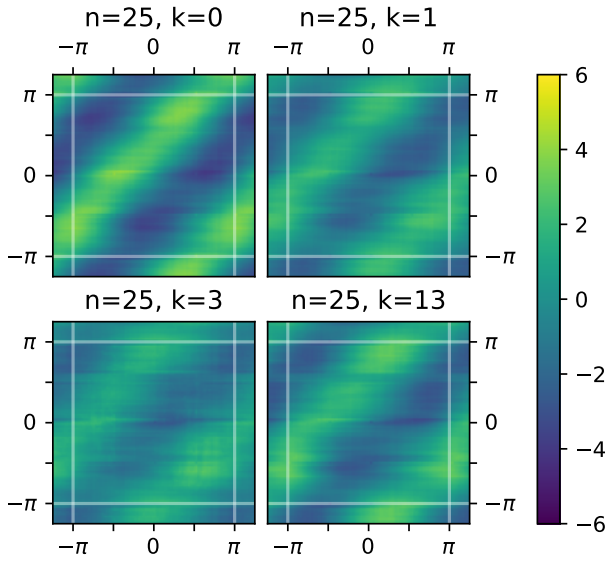

Figure 6.9: RMS ripple magnitude as shown in Figure 6.7 for four different 25-turn coils.

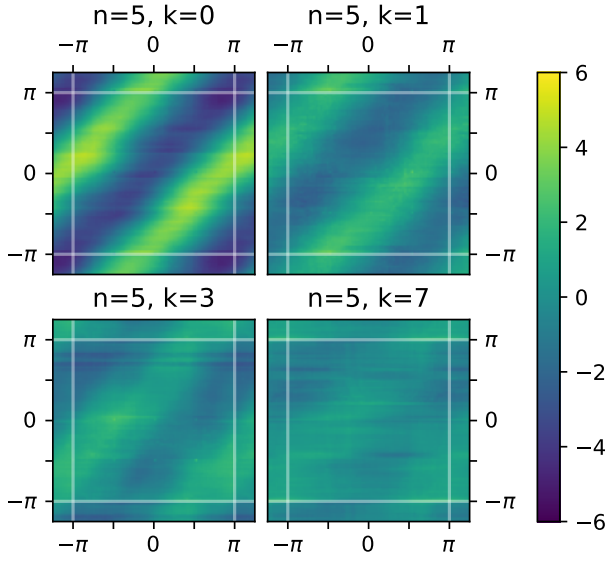

Figure 6.10: RMS ripple magnitude as shown in Figure 6.7 for four different 5-turn coils.

to 50 % improved self-resonant frequency as well as up to 6.5 % increased inductance compared to conventional two-layer planar spiral inductors.

We base our evaluation on laboratory measurements on a set of 39 sample inductors in total, including an automated, four-dimensional mapping of the coupling between a pair of identical inductors. We provide both an analytical description of twisted inductor construction as well as a set of Open-Source tools for their design in the supplementary material to this thesis.

Applied to an IHSM design, a wireless power transfer system using twised inductors to power the rotating mesh improves efficiency by reducing losses due to stray capacitance and reduces secondary-side ripple. The reduced secondary-side ripple allows the use of smaller filtering components, reducing board mass and mitigating heavy components as a possible fault location. Additionally, the reduced ripple allows the use of secondary-side voltage regulators with less voltage headroom, further reducing power transfer losses. By directly embedding twisted inductors into the IHSM's secondary side mesh monitoring PCB, construction is simplified. The resulting assembly is lighter and smaller, which reduces motor load and enables the implementation of compact IHSM meshes.

Web sources

[W87] *Gerbonara: Tools to Handle Gerber and Excellon Files in Python.* Version 1.4.0 (cit. on p. 172).

[W134] *KiCad EDA.* URL: `https://www.kicad.org/` (visited on 2024-12-03) (cit. on p. 172).

[W222] Juha Ruokolainen et al. *ElmerCSC/Elmerfem: Elmer 9.0.* Version release-9.0. Zenodo, 2023-05-03. DOI: `10.5281/ZENODO.7892181` (cit. on p. 172).

References

[16] Clifford W. Ashley and Geoffrey Budworth. *The Ashley Book of Knots: With Amendments.* Reprint. New York: Doubleday, 1993. 620 pp. ISBN: 978-0-385-04025-9 (cit. on p. 155).

[18] Charles Marfo Awuah et al. "Novel Coil Design and Analysis for High-Power Wireless Power Transfer with Enhanced Q-factor". In: *Scientific Reports* 13.1 (2023-03-14), p. 4187. DOI: 10.1038/s41598-023-31389-y (cit. on p. 157).

[21] T. Batra, E. Schaltz, and S. Ahn. "Effect of Ferrite Addition above the Base Ferrite on the Coupling Factor of Wireless Power Transfer for Vehicle Applications". In: *Journal of Applied Physics* 117.17 (2015-05-07), p. 17D517. DOI: 10.1063/1.4919039 (cit. on pp. 157, 162).

[53] Brian Curran et al. "Modeling and Characterization of PCB Coils for Inductive Wireless Charging". In: *Wireless Power Transfer* 2.2 (2015-09), pp. 127–133. DOI: 10.1017/wpt.2015.14 (cit. on pp. 157, 162).

[55] M. Danesh and J.R. Long. "Differentially Driven Symmetric Microstrip Inductors". In: *IEEE Transactions on Microwave Theory and Techniques* 50.1 (2002-01), pp. 332–341. DOI: 10.1109/22.981285 (cit. on p. 163).

[64] F. Eppen. "Anforderungen an Die Einzelteile Der Rundfunkempfänger; Gesichtspunkte Für Den Bau Der Geräte". In: *Die Wissenschaftlichen Grundlagen Des Rundfunkempfangs*. Ed. by K. W. Wagner. Berlin: Verlag von Julius Springer, 1927 (cit. on p. 165).

[67] Yuanshuang Fan et al. "A Simultaneous Wireless Power and Coil Inductance Insensitive Data Transfer System for Rotary Structures". In: *IEEE Transactions on Power Electronics* 39.5 (2024-05), pp. 6526–6536. DOI: 10.1109/TPEL.2024.3367295 (cit. on pp. 157, 159).

[68] Fritz Filbig. *Lehrbuch Der Hochfrequenztechnik*. Vol. 1. Akad. Verlag Becker & Erler, 1942 (cit. on p. 165).

[70] J. A. Fleming. *The Principles of Electric Wave Telegraphy and Telephony*. 2nd ed. Longmans, Green, and Co., 1910 (cit. on pp. 163, 164).

[87] *Gerbonara: Tools to Handle Gerber and Excellon Files in Python*. Version 1.4.0 (cit. on p. 172).

[89] Christophe Geuzaine and Jean-François Remacle. “Gmsh: A 3-D Finite Element Mesh Generator with Built-in Pre- and Post-processing Facilities”. In: *International Journal for Numerical Methods in Engineering* 79.11 (2009-09-10), pp. 1309–1331. DOI: 10.1002/nme.2579 (cit. on p. 172).

[112] Heng-Ming Hsu et al. “Analytical Design Algorithm of Planar Inductor Layout in CMOS Technology”. In: *IEEE Transactions on Electron Devices* 55.11 (2008-11), pp. 3208–3213. DOI: 10.1109/TED.2008.2004248 (cit. on p. 163).

[137] Paul-Eduard Klein. *Spulen Und Schwingungskreise*. Deutsche Radio-Bücherei 60. Deutsch-Literarisches Institut J. Schneider, 1941 (cit. on pp. 158, 165).

[146] Chi Kwan Lee, Y. P. Su, and S. Y. Ron Hui. “Printed Spiral Winding Inductor With Wide Frequency Bandwidth”. In: *IEEE Transactions on Power Electronics* 26.10 (2011-10), pp. 2936–2945. DOI: 10.1109/TPEL.2010.2076318 (cit. on pp. 164, 167).

[148] Woncheol Lee et al. “A Simple Wireless Power Charging Antenna System: Evaluation of Ferrite Sheet”. In: *IEEE Transactions on Magnetics* 53.7 (2017-07), pp. 1–5. DOI: 10.1109/TMAG.2017.2676099 (cit. on pp. 157, 162).

[151] Siqi Li and Chunting Chris Mi. “Wireless Power Transfer for Electric Vehicle Applications”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 3.1 (2015-03), pp. 4–17. DOI: 10.1109/JESTPE.2014.2319453 (cit. on pp. 157, 159, 161).

[152] Tao Li et al. “Wireless Power Transfer System for Long-term Sensor on Rotating Plane”. In: *2021 IEEE Industrial Electronics and Applications Conference (IEACon)*. 2021 IEEE Industrial Electronics and Applications Conference (IEACon). 2021-11, pp. 136–140. DOI: 10.1109/IEACon51066.2021.9654747 (cit. on p. 159).

[154] Ignacio Lope, Claudio Carretero, and Jesus Acero. “First Self-resonant Frequency of Power Inductors Based on Approximated Corrected Stray Capacitances”. In: *IET Power Electronics* 14.2 (2021-02), pp. 257–267. DOI: 10.1049/pel2.12030 (cit. on pp. 164, 167).

[155] Ignacio Lope et al. "Frequency-Dependent Resistance of Planar Coils in Printed Circuit Board With Litz Structure". In: *IEEE Transactions on Magnetics* 50.12 (2014-12), pp. 1–9. DOI: 10.1109/TMAG.2014.2337836 (cit. on p. 163).

[156] J.M. Lopez-Villegas et al. "Improvement of the Quality Factor of RF Integrated Inductors by Layout Optimization". In: *IEEE Transactions on Microwave Theory and Techniques* 48.1 (2000-01), pp. 76–83. DOI: 10.1109/22.817474 (cit. on p. 163).

[162] David Maier et al. "Contribution to the System Design of Contactless Energy Transfer Systems". In: *IEEE Transactions on Industry Applications* 55.1 (2019-01), pp. 316–326. DOI: 10.1109/TIA.2018.2866247 (cit. on p. 157).

[166] Peter Martin, Richard Horn, and Kobi Ben Atar. "A Multi-Turn Twisted Inductor for on-Chip Cross-Talk Reduction". In: *2016 IEEE International Conference on the Science of Electrical Engineering (ICSEE)*. 2016 IEEE International Conference on the Science of Electrical Engineering (ICSEE). 2016-11, pp. 1–5. DOI: 10.1109/ICSEE.2016.7806138 (cit. on p. 163).

[168] David Meeeker. *Finite Element Method Magnetics. User's Manual.* 2015-10-25. 161 pp. (cit. on p. 173).

[169] H. Meinke and F. W. Gundlach. *Taschenbuch Der Hochfrequenztechnik*. Springer-Verlag, 1956 (cit. on p. 165).

[175] S.S. Mohan et al. "Simple Accurate Expressions for Planar Spiral Inductances". In: *IEEE Journal of Solid-State Circuits* 34.10 (1999-10), pp. 1419–1424. DOI: 10.1109/4.792620 (cit. on p. 170).

[177] Julian Moore et al. "Applications of Wireless Power Transfer in Medicine: State-of-the-Art Reviews". In: *Annals of Biomedical Engineering* 47.1 (2019-01), pp. 22–38. DOI: 10.1007/s10439-018-02142-8 (cit. on pp. 157, 161).

[181] Xiaolin Mou, Oliver Groling, and Hongjian Sun. "Energy-Efficient and Adaptive Design for Wireless Power Transfer in Electric Vehicles". In: *IEEE Transactions on Industrial Electronics* 64.9 (2017-09), pp. 7250–7260. DOI: 10.1109/TIE.2017.2686299 (cit. on pp. 157, 159, 161).

[182] Xiaolin Mou and Hongjian Sun. “Wireless Power Transfer: Survey and Roadmap”. In: *2015 IEEE 81st Vehicular Technology Conference (VTC Spring)*. 2015 IEEE 81st Vehicular Technology Conference (VTC Spring). 2015-05, pp. 1–5. DOI: 10.1109/VTCSpring.2015.7146165 (cit. on pp. 157, 161).

[183] U. Muehlmann, M. Gebhart, and M. Wobak. “Mutual Coupling Modeling of NFC Antennas by Using Open-Source CAD/FEM Tools”. In: *2012 IEEE International Conference on RFID-Technologies and Applications (RFID-TA)*. 2012 IEEE International Conference on RFID-Technologies and Applications (RFID-TA). 2012-11, pp. 393–397. DOI: 10.1109/RFID-TA.2012.6404553 (cit. on p. 162).

[184] 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 (1989-03), pp. 287–292. DOI: 10.1007/BF02657420 (cit. on pp. 136, 163).

[186] Eugen Nesper. *Handbuch Der Drahtlosen Telegraphie Und Telphonie*. Vol. 2. 2 vols. Julius Springer, 1921 (cit. on p. 163).

[188] Minh Huy Nguyen and Handy Fortin Blanchette. “A Review and Comparison of Solid, Multi-Strands and Litz Style PCB Winding”. In: *Electronics* 9.8 (2020-08-16), p. 1324. DOI: 10.3390/electronics9081324 (cit. on p. 163).

[190] Shunsaku Nomoto et al. “Splitting Conductors of Coils on PCB for AC-resistance Reduction”. In: *2024 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2024 IEEE Applied Power Electronics Conference and Exposition (APEC). 2024-02, pp. 3204–3209. DOI: 10.1109/APEC48139.2024.10509283 (cit. on p. 163).

[192] Heinrich Nottebrock. *Spulen*. Vol. 3. Bauelemente Der Nachrichtentechnik. Schiele & Schön, 1950 (cit. on p. 165).

[211] William Querfurth. *Coil Winding: A Description of Coil Winding Procedures, Winding Machines and Associated Equipment*. G. Stevens Mfg. Company, 1954 (cit. on p. 165).

[222] Juha Ruokolainen et al. *ElmerCSC/Elmerfem: Elmer 9.0*. Version release-9.0. Zenodo, 2023-05-03. DOI: 10.5281/ZENODO.7892181 (cit. on p. 172).

[224] Hermann Saacke. *Radiotechnik III: Die Empfänger*. Vol. 3. Sammlung Göschen. Walter de Gruyter & Co., 1926 (cit. on p. 158).

[237] Kai Song et al. "A Rotation-Lightweight Wireless Power Transfer System for Solar Wing Driving". In: *IEEE Transactions on Power Electronics* 34.9 (2019-09), pp. 8816–8830. DOI: 10.1109/TPEL.2018.2886910 (cit. on p. 159).

[239] Ole Christian Spro, Frank Mauseth, and Dimosthenis Peftitsis. "High-Voltage Insulation Design of Coreless, Planar PCB Transformers for Multi-MHz Power Supplies". In: *IEEE Transactions on Power Electronics* 36.8 (2021-08), pp. 8658–8671. DOI: 10.1109/TPEL.2021.3049353 (cit. on p. 167).

[242] M. J. O. Strutt. *Verstärker Und Empfänger*. 2nd ed. Vol. 4. Lehrbuch Der Drahtlosen Nachrichtentechnik. Springer-Verlag, 1951 (cit. on p. 165).

[267] Longyang Wang et al. "Coaxial Nested Couplers-Based Offset-Tolerance Rotary Wireless Power Transfer Systems for Electric Excitation Motors". In: *IEEE Access* 8 (2020), pp. 44913–44923. DOI: 10.1109/ACCESS.2020.2978130 (cit. on p. 159).

[268] Qiyue Wang, De'an Wang, and Jiantao Zhang. "A Novel Rotating Wireless Power Transfer System for Slipring with Redundancy Enhancement Characteristics". In: *Sustainability* 16.13 (13 2024-01), p. 5628. DOI: 10.3390/su16135628 (cit. on pp. 159, 162).

[276] Heinrich Wigge. *Rundfunktechnisches Handbuch*. 2nd ed. Vol. 1. Verlag von M. Krayn, 1930 (cit. on p. 165).

[278] Kun Xia et al. "A Rotary Wireless Power Transfer System With Rail-Type Coupling Structure". In: *IEEE Access* 12 (2024), pp. 63967–63975. DOI: 10.1109/ACCESS.2024.3393943 (cit. on p. 159).

[279] Zhengchao Yan et al. "Free-Rotation Wireless Power Transfer System Based on Composite Anti-Misalignment Method for AUVs". In: *IEEE Transactions on Power Electronics* 38.4 (2023-04), pp. 4262–4266. DOI: 10.1109/TPEL.2023.3238066 (cit. on p. 159).

[282] Yiming Zhang et al. “An Improved Compensation Method Reducing Displacement Current Loss for Multilayer Coils in IPT System”. In: *IEEE Transactions on Power Electronics* 40.1 (2025-01), pp. 87–91. DOI: 10.1109/TPEL.2024.3462669 (cit. on p. 174).

[283] Zeheng Zhang et al. “A Dynamic Wireless Power Transfer System Using DC-Controlled Variable Inductor for Segment Transmitter Automatic Switching”. In: *IEEE Transactions on Power Electronics* 40.1 (2025-01), pp. 23–27. DOI: 10.1109/TPEL.2024.3426100 (cit. on p. 162).

[284] Zeheng Zhang et al. “Wireless Sensor Power Supply for Rotating Shaft Using DC-Side Diode Array With Stable Output”. In: *IEEE Transactions on Power Electronics* 39.12 (2024-12), pp. 15414–15419. DOI: 10.1109/TPEL.2024.3439718 (cit. on p. 159).

[285] Zhen Zhang et al. “Wireless Power Transfer—An Overview”. In: *IEEE Transactions on Industrial Electronics* 66.2 (2019-02), pp. 1044–1058. DOI: 10.1109/TIE.2018.2835378 (cit. on pp. 157, 161).

[288] G. Zickner. “Spulen”. In: *Taschenbuch Der Drahtlosen Telegraphie Und Telephonie*. Ed. by Fritz Banneitz. Verlag von Julius Springer, 1927 (cit. on p. 165).

Chapter 7

Case Study: Physical Security in Quantum Key Distribution

> One should always assume that people willing to break a system are also willing to use significantly more resources doing so than legitimate users are willing to spend routinely.
>
> – *Russell Impagliazzo [119]*

Contents

Quantum Computing promises efficient solutions to a number of widely used cryptographic computational problems. As a countermeasure, new *post-quantum* cryptosystems have been developed that are not susceptible to known quantum or classical attacks. However, a limitation of these cryptosystems is that they still rely on hardness assumptions that cannot be proven—and it cannot be ruled out that attacks on these cryptosystems could be found in the future. In fact, a variant of one of the early contenders for post-quantum cryptography, Supersingular Isogeny Diffie-Hellman Key Exchange (SIKE) has unexpectedly been broken in 2022 [40], a decade after its development, highlighting the risk inherent in these new cryptosystems.

Quantum Key Distribution (QKD) provides an alternative to key exchange protocols based on cryptographic hardness assumptions. QKD provides a primitive similar to Diffie-Hellman key exchange, establishing a secret key between two parties that are only connected through an untrusted channel. In contrast to classical cryptographic protocols, the security of QKD is based on quantum-physical laws of nature, and assuming a correct technical realization, QKD can provide information-theoretic security.

QKD suffers from a severe range limitation stemming from loss in optical fibers. Since QKD relies on the quantum properties of single photons, QKD signals inherently cannot be amplified. While classical optical networking signals can be efficiently amplified using optical amplifiers, to a QKD signal such amplification would constitute a measurement, which destroys the signal's quantum information. As a consequence of this, the range of a QKD link is limited to the span that can be achieved with a single, uninterrupted fiber at an acceptable loss. In practice, this is commonly in the range of 100 km to 200 km with key exchange rates falling sharply with longer distance.

The only technique for range extension that is currently feasible is to *relay* the QKD signal with a receiver and a transmitter coupled back-to-back. This practical construction however creates another hard challenge: Since only the QKD system's photonic signal is secured by the systems' quantum security guarantees, such relays must be physically trusted as they effectively handle secret key bits in plaintext. Achieving this physical security in a large-scale QKD network is difficult due to the remote location of some relays, the QKD nodes' physical size, and their power and cooling requirements, and their need for multiple fiber-optic connections to the outside world. In classical computing, such challenges are often approached us-

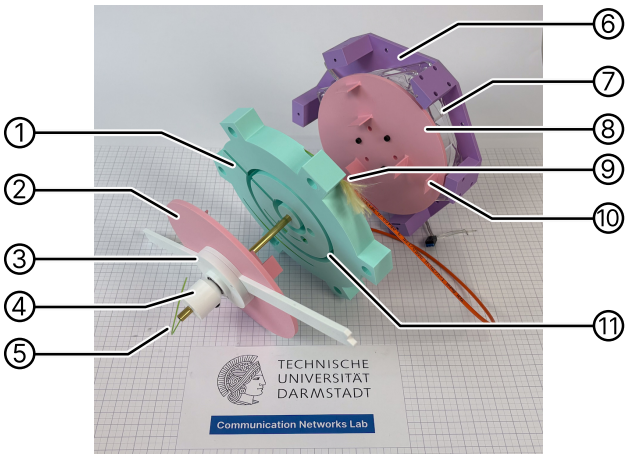

Figure 7.1: Photo of our mechanical prototype. 1 - Bracket connecting payload and shaft with hidden spiral conduit for optical fibers. 2 - Upper tamper sensing mesh PCB. 3 - Outer IHSM tamper sensing mesh cage. 4 - IHSM tamper sensing mesh cage bearing. 5 - Fiber exiting hollow shaft. 6 - Lower bracket holding secondary tamper sensing mesh drive motor. 7 - Cooling fan used as secondary tamper sensing mesh drive motor. 8 - Secondary tamper sensing mesh PCB shielding bottom of bracket 1. 9 - Fiber exiting hidden spiral conduit in bracket 1. 10 - Interleaving tabs sticking out from tamper sensing PCBs, creating a serpentine structure. Distance from tab end to opposing PCB 2 is 3.4 mm of space in 11 - Channels for tabs 10 in bracket 1.

ing Hardware Security Modules (HSMs) that have tamper sensors that will destroy the HSM's contents when tampering is detected, but conventional HSM technology cannot be adapted to the requirements of a QKD system.

In this chapter, we present several designs and a mechanical prototype adapting the Inertial Hardware Security Module (IHSM) concept first proposed by Götte and Scheuermann [94] to a QKD relay node. IHSMs replace the tamper sensing security mesh foil that is wrapped around the payload in conventional HSMs by a tamper-sensing cage made from conventional circuit board material by spinning this cage at a high speed. On its own, circuit board material provides lower tamper security than the tamper sensing foils made using bespoke manufacturing processes that are used in conventional HSMs. IHSMs solve this problem by spinning the tamper sensing cage at high speed while continuously verifying this rotation using an accelerometer placed on the cage. IHSMs achieve a similar security level to conventional HSMs using only inexpensive, commodity components and no specialty manufacturing processes. In contrast to conventional HSMs, IHSMs are a natural fit for the high power and size requirements of a QKD

node. However, they suffer from the problem of how to optically connect the (stationary) QKD relay payload protected inside the IHSM's spinning tamper sensing cage to the outside world without creating a security vulnerability. While fibers can easily be fed through the shaft of the spinning cage, an attacker could feed an attack tool through the same opening. In this chapter, we propose a family of mechanical designs that use a secondary rotating tamper sensing mesh at the entry point of the shaft to protect a fiber-optical passthrough while observing the fiber's bending radius limitations. Figure 7.1 shows a photo of our mechanical prototype. Our prototype would require an attacker to feed an attack tool around multiple sharp bends, with only 3.4 mm of space available at the narrowest points. In our prototype, the smallest bend radius encountered by the fiber is 15 mm. We experimentally measured the optical loss added by our prototype compared to a straight fiber to be below our measurement floor of 0.25 dB.

This chapter is organized as follows. In Section 1, we give an introduction into Quantum Key Distribution and its practical realization. In Section 2, we provide an overview of related academic work. In Section 3, we introduce three variants of our optical passthrough design that lie along different points of the security/complexity spectrum. In Section 4 we discuss attacks on our design before concluding with an outlook of future research directions in Section 5.

1 QKD Fundamentals

In principle, QKD is a specialized form of photonic quantum computing. The underlying approach in QKD is that two parties exchange quantum states, then perform experiments on these quantum states to produce partially correlated randomness. This correlated randomness is then refined into identical secrets on both ends by running an error correction process known as *information reconciliation* using a classical channel for communication. After this process, an attacker may still possess partial information about the shared secret. To dilute this information, in a step named privacy amplification, a randomness extractor such as a information-theoretic hash function is used to create a new, shorter secret over which the attacker possesses effectively no information.

1.1 Range in QKD

Regardless of the particular QKD protocol used, common to all QKD protocols, quantum states must be exchanged between parties. While quantum computers are built from a wide variety of quantum states from trapped ions through superconducting states up to spin states, all QKD protocols are based on photonic states since they are the only ones that can easily be transferred across long distances through optical fiber. Even so, QKD protocols face a steep trade-off between speed of key generation—called *secret key rate*—and distance since quantum states cannot be amplified. In literature on long-range QKD, secret key rates as low as 10 milli-bits per second are routinely published [269] since they already promise a benefit in a hypothetical scenario in which symmetric cryptography cannot yet be efficiently attacked using Grover's algorithm, but all asymmetric cryptography has fallen to quantum algorithms like variants of Shor's algorithm.

1.2 Loss in optical fibers

When transmitted over a fiber, there are multiple effects that degrade the quantum-optical signal of a QKD system, which are collectively referred to as *loss*. We can coarsely classify these degrading effects into two categories: *decoherence*, and *attenuation*. Decoherence effects result in the quantum state being changed in transit, which depending on the QKD implementation may mean destroying information contained within the state such as by disturbing the pulse's polarization, or destruction of entanglement between the in-flight state and another local state.

Decoherence effects are less relevant for the distance limitation, and mostly limit which fiber-optic technologies can be utilized in the first place. Due to decoherence, QKD systems usually use Single-Mode (SM) fiber over Multi-Mode (MM) fiber [11], and decoherence makes it more difficult to utilize Wavelength Division Multiplexing (xWDM) to send multiple either quantum or classical optical signals through a single fiber.

In practice, attenuation is the primary factor limiting the length of an individual fiber run in QKD. Even modern, ultra-low loss optical fiber has an attenuation in the order of 0.15 $\frac{\text{dB}}{\text{km}}$, resulting in a loss of half the signal's power, equivalent to half of all QKD pulses, in just 20 km. Since these losses compound exponentially with longer reach, after only 200 km only one in a thousand photons entering the fiber will exit it at the other end [44].

1.3 Relaying

A consequence of this range limitation is that at useful bit rates, QKD links can only be realized up to distances in the order of 200 km. There are some QKD protocols that can be used to effectively double the range of a QKD link by placing an untrusted node in the middle of the link, but further extension would require either a trusted relay or a complex relay operating on the quantum states. As of now, such quantum relays are not practical leaving only the trusted relay route for achieving useful secret key rates across distances longer than a few hundred kilometers.

If we imagine a continental-scale network of QKD systems with fibers spanning tens of thousands of kilometers, it is easy to see why the physical security of its relay nodes is such a concern in QKD setups. Such a network would need between hundreds and throusands of relay nodes. Making things worse, these relay nodes would have to be spread evenly across thousands of kilometers of optical links, with many ending up in isolated places in the field, away from datacenters and other well-protected technical infrastructure. Since the compromise of any one QKD relay could be enough for an attacker to carry out an on-path attack, protecting thousands of small relay installations located in equipment sheds spread across sparsely populated areas against adversaries with advanced physical attack capabilites becomes a daunting task. Effectively, each quantum relay has to be made into a hardware security module including advanced active tamper sensing.

2 Related Work

2.1 Long-range QKD

Cao et al. [38] give a comprehensive overview of large-scale QKD networking. Lella and Schmid [150] analyze security threats in quantum key distribution networks and point out that achieving the information-theoretic security that QKD is often cited for providing is difficult to achieve in practice since currently, protocols based on cryptographic computational hardness assumptions cannot be avoided in a practical implementation. Yang, Zhang, and Su [280] approach key routing in a hypothetical quantum key distribution network and provide a solution based on measurements of each node's local secret key buffer.

Cao et al. [37] discuss hybrid QKD networks that employ both physically trusted and untrusted nodes by applying a technique such as Measurement-

Device Independent QKD (MDI-QKD) that enables one end of the QKD link to be untrusted. MDI-QKD can effectively double the reach of a trusted QKD link by placing an untrusted relay node in the middle. They present a precise problem formulation and introduce an algorithm for the optimization of deployment cost of a hybrid QKD network.

2.2 Customizable tamper sensing HSMs

Immler et al. [117] introduce a HSM concept that utilizes a tamper-sensing mesh made from a lithographically patterned metallized polyimide foil. They pattern a grid of fine capacitive electrodes onto the foil, and demonstrate a simple multi-channel readout circuit that is capable of distinguishing changes in capacitance between electrodes down to the femto-Farad range. In contrast to conventional HSMs that require a continuous power supply to their tamper-sensing subsystem, their design introduces sufficient measurement fidelity that the tamper-sensing mesh foil can be viewed as a Physically Uncloneable Function (PUF) by demonstrating stability and statistical properties of its PUF response.

Later publications on their design expand upon the concept, but fundamentally, their design is limited in size by manufacturing limitations in the size of its tamper-sensing foil, as well as the poor scalability of the designs frontend architecture, which requires a separate charge amplifier for each electrode pair [78, 79, 80, 197]. Applying their approach to a QKD relay would be difficult as it would require not just miniaturizing the QKD relay to the size of a smartphone, but it would also require the development of a secure fiber passthrough specific to their design and other systems using a folded tamper-sensing mesh foil. Conventionally, electrical pass-throughs in such foils are made by folding the mesh and a Flat Flexible Cable (FFC) multiple times. Due to their required beding radius, alternative solutions would have to be found for a fiber-optic pass-through.

2.3 Inertial Hardware Security Modules

As of now, QKD nodes are large, rack-mount devices. While miniaturization is ongoing, the processing requirements of such systems alone exceed the capabilities of conventional HSMs. With a conventional HSM, protecting an entire QKD relay consisting of two link endpoints and their associated processing systems would be infeasible due to their size and power dissipation.

One of the core challenges in the design of active tamper sensors for HSMs is protecting the device against drilling attacks. In a drilling attack, an attacker accesses internal circuitry of the HSM by drilling a hole, allowing a probe to pass through. In HSMs, drilling attacks are commonly monitored by enveloping the payload in a security mesh, i.e. a foil covered with intentionally fragile conductive traces. The idea is that drilling into the device from any angle will damage the conductive traces on this foil, which can easily be electrically detected by the payload, allowing it to destroy all secrets before any probe can reach it.

In practice, manufacturing this conductive foil is difficult. Standard flexible circuit processes such as lithographic polyimide/copper Flexible Printed Circuits (FPCs) are sometimes used, but their security is limited since they are easy to manipulate using standard Printed Circuit Board (PCB) rework techniques. More exotic processes industrially used for low-cost keyboard and key pad production using screen-printed silver or carbon conductive inks on a polyester substrate are also used, but are limited by a coarse structure size.

The area of foil-based security meshes is primarily limited by the difficulty of manufacturing large foils without defects. Not only does total defect rate rise with area, commercial PCB or FPC manufacturing processes have a panel size usually in the order of 500 mm to 800 mm side length that cannot be exceeded.

In contrast to conventional HSMs using mesh foils, IHSMs approach envelope tamper sensing by encasing the payload in a mesh cage made from low-cost PCBs, then rotating this cage at high speed to simultaneously cover all angles, and prevent manipulation of the mesh. To prevent an attacker from slowing down the rotating mesh cage, an accelerometer is placed on the rotating mesh that monitors rotation by measuring centrifugal acceleration.

The main issue in IHSM construction is the construction of the passthrough providing electrical connections between the payload and the outside world. In conventional HSMs that use tamper sensing mesh foils, this passthrough is realized by folding the mesh foil and a Flexible Flat Cable (FFC) in several layers such that there is no straight path that a probe could be inserted through. In IHSMs, electrical connections are passed through a hollow shaft on one end of the mesh cage. Similar to the serpentine folds between mesh foil and FFC in conventional HSMs, in IHSMs complex geometry can be realized by placing a secondary rotating mesh on the inside

of the primary mesh, covering the point where the shaft goes through the primary mesh.

Where in conventional HSMs covering larger areas with a patchwork of smaller mesh foils creates the difficulty of creating secure seams between the foils, in IHSMs, multiple PCB meshes can easily be joint into a larger mesh by simply overlapping them, since the mesh's rotation makes any attack on such a joint exceedingly difficult.

3 Multi-fiber passthrough with active secondary mesh

Since IHSMs are particularly suited to large payloads, fitting the components of a QKD node inside one is straightforward. However, QKD links have one unique requirement: Many systems require several physical fibers for each QKD link. Often, in addition to a fiber for classical communication, one fiber is needed to transmit a reference clock to the other end of the link, and another fiber is needed for the quantum channel. With a QKD relay needing at least two links, this results in at least five fibers assuming all classical networking can be multiplexed on a single fiber.

Fiber pigtails have an outer diameter of usually about 1 mm, so this amount of fibers can be fed through an IHSM's axis of rotation without increasing its shaft diameter and reducing its security. The mechanical challenge in such a multi-fiber signal and data feedthrough is to observe the fiber's minimum bending radius, which for common fibers is usually in the range of 5 mm to 15 mm [W74, W210, W50].

3.1 Multi-fiber passthrough design

To approach the security of the data and power connections passing through the IHSM's unprotected shaft, Götte and Scheuermann [94] list some shielding methods that use an independently rotating secondary tamper sensing mesh on the inside of the primary mesh, located right next to the primary mesh's axis opening. This secondary mesh makes accessing the payload using probes inserted through the shaft much more difficult. Götte and Scheuermann [94] only present conceptual drawings of these schemes, and focus on electrical signals. In this chapter, building on these concepts, we present mechanical designs of three variations of a fiber passthrough for IHSMs that are adapted to the limited bending radius of optical fiber: A

simple disc cover, offset labyrinth meshes, and interlocking gear meshes. We present a mechanical prototype of our offset labyrinth mesh design.

3.2 Simple disc cover

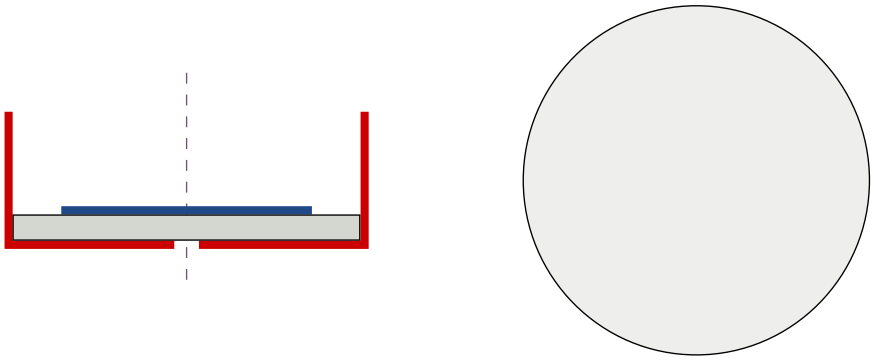

Figure 7.2: Coaxial disc mesh schema, cross-section and top-down views. The outer mesh is shown in red, and the inner mesh in blue. The dashed line indicates the two meshes' shared axis of rotation. The gray areas indicate the shape of the volume that remains undisturbed by the mesh, and that is available for structural support and cable routing.

While IHSMs excel at protecting large payload volumes, even a zero-payload IHSM that has been shrunk to a single, disc-shaped PCB is still useful because we can delegate key management functionality to the mesh monitoring circuit's microcontroller—or a separate processor sitting next to it—on the rotating mesh PCB, yielding a solution close in both its cryptographic capabilities and its security level to commercial traditional HSMs, and exceeding those of a smartcard. In the following paragraphs, we will show how we can deploy the same single-board IHSM (SB-IHSM) as a mitigation for through-axis attacks, exploiting its mechanical shape and its simple, low-cost implementation.

By placing an adapted single-board IHSM close to the primary mesh's axis opening as shown in Figure 7.2, an attacker is forced to either first circumvent or at least dislodge the single-board IHSM through the primary mesh's axis opening without disturbing either mesh to gain direct access to the payload behind it, or to conduct their attack through the keyhole-sized opening in the primary mesh while bending their tool by approximately 90° at least twice, once to avoid the SB-IHSM mesh, and once more to re-orient the tool towards the payload. The distance between the inside of the primary mesh and the SB-IHSM is limited by the tolerance in me-

chanical alignment between the two axes of rotation, by the space necessary for a sufficiently stable mount of the payload cage to the hollow shaft, and by the minimum bend radius of the power and data wiring that needs to pass through the shaft. Power and electrical data signals can be supplied through flexible flat cables that can be bent in sharp corners without issue. In QKD applications, the fibers' minimum bend radius is the largest contributing factor. The optical loss of a fiber rises sharply with decreasing bend radius[1]Note that the issue here is not that the glass core of the fiber would degrade or break, as one might intuitively assume. Being only a few dozen micrometers in diameter, an optical fiber's core is remarkably flexible. Instead, the issue is that both multi-mode as well as single-mode fibers are optical waveguides. Bending them distorts the electromagnetic field inside the waveguide, and allows some small portion of it to escape from the fiber's core, leading to loss in the form of both attenuation and dispersion [229]. With QKD being especially sensitive to even small amounts of loss, care has to be taken to maximize the bend radius of the fiber optic connections. A common specification of minimum bend radius in telecom single-mode fibers taking into account not just optical loss but also the mechanical stability of the fiber's polymer coating is 10× the coated fiber's diameter [W74, W210, W50], which equates to 9 mm for common 0.9 mm fiber pigtails, corresponding to approximately 1 dB of loss in the 1550 nm band [229]. Based on these specifications and on a conservative estimate of 2.5 mm for the vertical mesh clearance, we arrive at a minimum inter-mesh spacing of approximately 11 mm when using minimal overlap between tab heights.

3.3 Coaxial labyrinth meshes

In QKD applications, the simple disc cover design shown above has two main limitations. First, the distance between the primary and secondary meshes' tab rings must be large enough to allow for the fibers' minimum bend radius, resulting in more than 10 mm of space available to an attacker. Second, the attacker only has to bend their tool in a plane to reach the payload.

To increase the difficulty of inserting a long and flexible tool through the axis shield, the shape of the interface layer between the two meshes can be made more complex. Introducing small mesh *tabs* that stick out into

[1] .

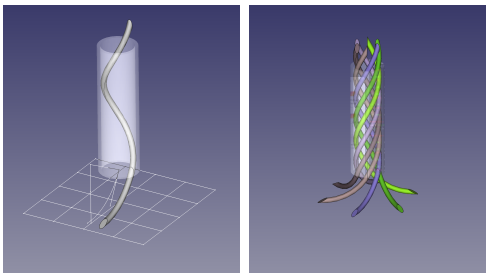

(a) Single fiber (b) Fiber bundle

Figure 7.3: Minimum mesh spacing can be reduced by coiling the fibers inside of the shaft tube. The coiled fibers enter the inter-mesh space at an angle equal to the helix lead angle. Shown here is a 6 mm outer diameter tube with a 0.5 mm wall thickness and 6 fibers with 0.9 mm outer diameter coiled to a constant bend radius of 9 mm. The lead angle of the helix is 61.5°. The resulting inter-mesh spacing is 5.16 mm.

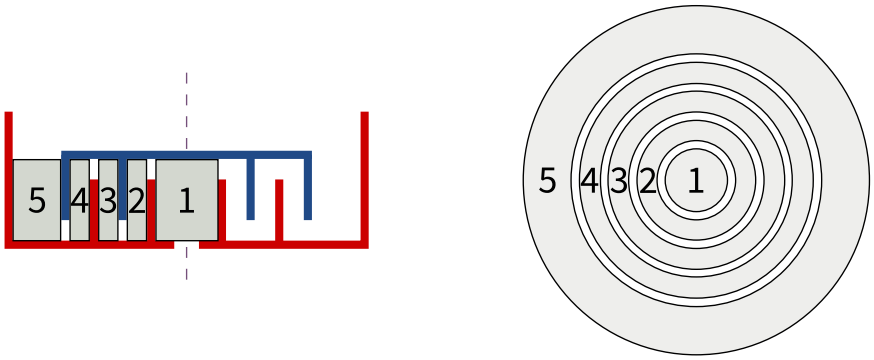

Figure 7.4: Coaxial labyrinth mesh schema, cross-section and top-down views.

the inter-mesh space from both meshes creates a labyrinth-like structure between the axis opening and the IHSM's inside. Structural support and cables can easily pass this structure in a series of 90° bends, while inserting a probe avoiding both meshes would not be feasible as the probe would have to perform a series of sharp bends. The type of manipulator that would be necessary for the placement of a probe in this system is conceptually similar to snake-like robots used in minimally invasive surgery, but state-of-the-art systems from this area are both too thick and don't have enough joints to fit even simple labyrinth layouts [243, 230, 135, 108]. For instance, if we assume 3 mm material thickness on the radial bracket connecting the shaft with the secondary mesh's mounting frame along with 10 mm of mesh tab overlap, 1.5 mm of clearance between radial bracket and each of the two meshes, and an inter-mesh spacing from one tab ring to the next equal

to the radial brackets' material thickness of 4 mm plus the clearance from bracket to mesh, we arrive at a meander 6 mm in width completing four 180° turns within less than 40 mm of radial distance.

While long and narrow tabs are desirable for mesh security as they limit the size and mobility of an attacker's probe, in QKD application, the need for fiber optic passthrough is the limiting factor. The obvious solution of passing through the fibers in a series of in-plane S-bends requires a coarse tab spacing due to the fibers' large minimum bend radius. However, we can apply the approach we proposed above for the shaft entrance here, too, and thread the fibers between the meshes by helically coiling them, increasing the fibers' bend radius to one half of the distance between both mesh discs minus the fibers' diameter and clearances. When the resulting useable part of the distance is larger than twice the bend radius, the minimum tab spacing is only limited by the fiber's diameter and the stability of the star bracket. When the discs are placed closer, and a larger pitch is necssary, the resulting pitch of the helix determines the minimum tab spacing.

Designing a labyrinth mesh for intrusion prevention is similar to the design of the shape of the jamb of a safe door or of a high end apartment door. In these, the objective is to prevent would-be burglars from inserting opening tools through the space between the closed door and its jamb and attacking the door's interior handle or locking mechanism, not unlike an IHSM's defense against electrical or electromagnetic probes. The one difference between these doors and what we can do in IHSMs is that these doors are limited to outwards-facing steps because they must be opened and closed. In IHSM labyrinth meshes, we can use both outwards-facing and inwards-facing steps.

Concentric labyrinth meshes allow for a range configurations. The pitch from one mesh tab to the next is the sum of the required width of the inter-mesh space and the safety margin needed betwween any cables or the inter-mesh bracket and the tabs. When the mesh is constructed using rigid PCB tabs that are inserted as-is, without bending them, and when all tabs have the same width and thickness, the radial width of the swept area decreases from tab to tab going outwards. A consequence of this is that when the design target are constant width inter-mesh spaces, the tabs' pitch decreases going outwards.

The safety margin required to avoid collisions between the meshes and the stator can be kept low for the primary mesh because this mesh has high-

quality bearings on both ends, leading to good axis alignment. In contrast, for the secondary mesh, margins have to be included if the mesh is driven by a cooling fan motor, as the bearings in such fans are not very precise, resulting in misalignment increasing with radius.

3.4 Offset labyrinth meshes

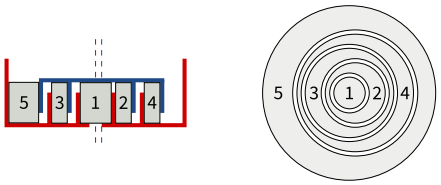

Figure 7.5: Offset labyrinth mesh schema, cross-section and top-down views. The two dashed lines indicate the two meshes' offset axes of rotation, shifted in x direction in both views.

Concentric labyrinth meshes improve upon simple disc meshes in security, but they have two remaining weaknesses. One is that in a concentric labyrinth mesh, the part of the inner mesh at the axis is easily accessible through the opening in the outer mesh. As the axis of rotation is the most vulnerable spot in a mesh because the tangential velocity of the mesh is lowest close to the axis, tampering can be made more difficult by placing the axis of rotation of the inner mesh not concentric with that of the outer mesh, but at a radial *offset*.

A consequence of placing the axis of the inner mesh at an offset is that the inter-mesh rings formed by the tabs of the two meshes now no longer form a set of concentric rings, but a set of nested non-concentric annulus shapes whose narrow and wide sides alternate along the direction of the offset. We will show below how an optical fiber can still be wound through this complex inter-mesh space without much trouble through a variation of the helical spiral trick from above to avoid the annular rings' narrow sections. At the same time, the alternating narrow sections of the annular rings make it more difficult to feed through the type of surgical robot we cited above, whose joints are designed for in-plane operation for most of the manipulator, starting from the high-flexibility joints close to its end and down the neck. In this section, we will show a design and a mechanical prototype of an offset labyrinth mesh design that improves on a concentric labyrinth mesh on both the shielding of the secondary mesh axis and the feasibility of an attack with a surgical robot without increasing mechanical

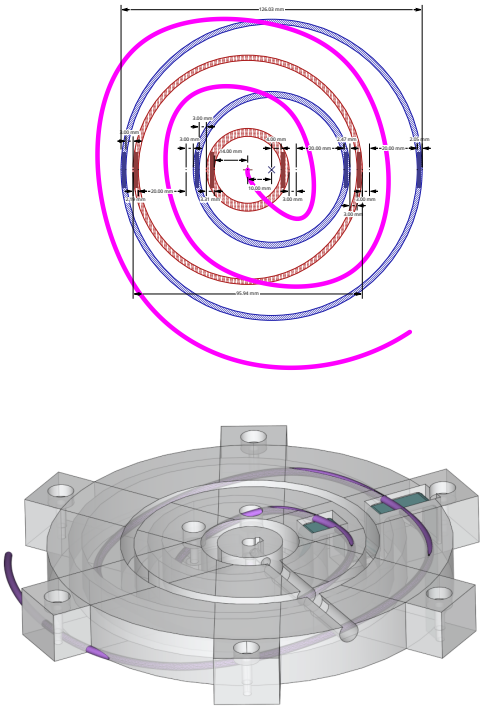

Figure 7.6: Offset labyrinth mesh schema with fiber layout

complexity compared to a concentric design. In addition, we show a fiber feedthrough that improves on the simple helical feedthrough we introduced above.

Our offset labyrinth mesh design combines an offset of the secondary mesh's axis of rotation with the labyrinth mesh approach from the previous section, creating wide and narrow inter-mesh spaces on alternating sides of the offset direction as shown in in Figure 7.5. Structural support is provided using a CNC machined or 3D printed part, which also serves as a conduit for electrical connections from the shaft to the payload using Flexible Flat Cable (FFC). While the FFC can easily conform to the offset labyrinth's sharp corners, an optical fiber can not. Thus, instead of passing it straight through the labyrinth, the payload's fiber optic connections are passed through the labyrinth in a three-dimensional spiral shape, avoiding the meshes while simultaneously maximizing the fibers' bend radii.

3.5 Experimental Validation

To prove the mechanical viability of the offset labyrinth mesh concept, we created a mechanical prototype of one such mesh. Figure 7.6 shows the

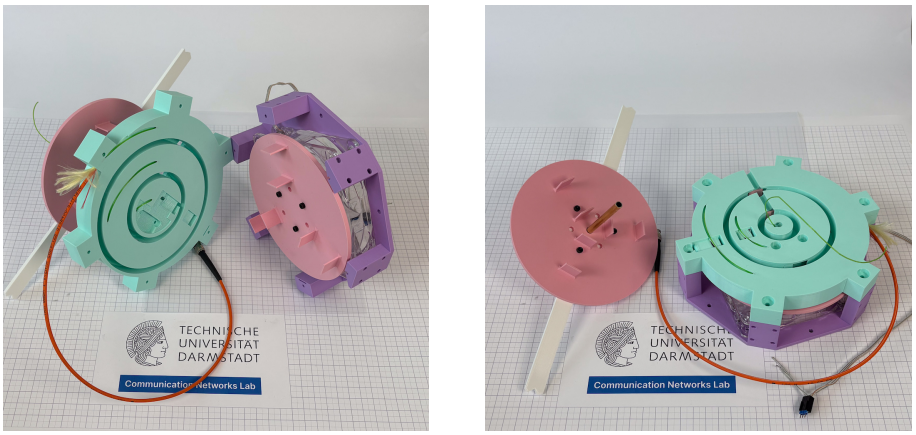

Figure 7.7: A disassembled view of our optical passthrough mechanical prototype. The fiber is passed through from the shaft going through the IHSM's primary tamper sensing mesh cage to the outside into the interior of the IHSM through the green bracket. A secondary tamper sensing mesh is located on the inside of the shaft interface and driven separately. In this prototype, the secondary mesh is driven by a cooling fan. Both independently rotating meshes have tabs that extend into the bracket such that they do not interfere, but reduce the space available to an attacker. The HSM's primary mesh cage is partially shown in white.

proportions of the meshes' tabs along with the resulting tab rings and a 2D projection of our chosen fiber layout. The fiber is laid out in such a way that it crosses each tab ring at opposite sides, and traverses the vertical distance in the larger part of the inter-mesh space. Figure 7.7 shows an exploded view of our mechanical prototype.

We threaded a standard 50 µm/125 µm fiber through the bracket, spliced it to a connector pigtail at the remote end, and measured its loss using a NK4000D handheld OTDR/OPM manufactured by Qingdao Novker Communication Ltd. Comparing measurements of loss between a coiled fiber and a fiber fed through the bracket resulted in a difference below the measurement floor of approximately 0.25 dB.

3.6 Interlocking gear meshes

The offset labyrinth design already achieves a high level of security through its complex passthrough shape, but only small offset distances are feasible since large offsets quickly lead to impractically large mesh sizes. Where the pitch from one tab ring to the next is roughly constant in concentric labyrinth meshes, and determined only by clearances and the amount of inter-mesh space necessary for power and data feedthroughs as well as me-

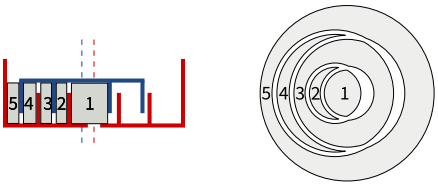

Figure 7.8: Offset gear labyrinth mesh schema, cross-section and top-down views. In this example, the axis is shifted by about twice the offset from the previous offset labyrinth mesh schema in Figure 7.5.

chanical stability. In offset meshes, on the other hand, this pitch increases by the offset distance. Even for a small offset this quickly adds up to an unwieldy total mesh size.

In this section, we conceptually introduce a solution to this problem that allows for larger offsets using a design where the two meshes interlock like gears. This does mean that the two meshes' rotation must be synchronized, but it increases the design space of offset labyrinth meshes. For instance, in a gear setup, the wide sides of the inter-mesh zones can be aligned to lie on the same side, so fiber passthrough can be realized more easily even without the need to spiral the fiber around the axes of rotation.

3.7 Mesh synchronization

For geared meshes to work, both speed and phase of the rotation of the two meshes must be synchronized to a small error. In this setup, the mesh tabs act like gear teeth. Depending on the ratio between both meshes' tap counts, the two meshes do not have to rotate at the same rate of rotation and harmonic ratios are possible. Additionally, unlike actual gears which need to constantly maintain an area of contact, both co-rotating and counter-rotating setups are possible.

4 Physical attacks and countermeasures

In this section we will consider possible ways to attack an IHSM-secured QKD relay, as well as potential countermeasures.

4.1 Attacks on the IHSM mesh

There are two ways an attacker could attack the mesh itself if an adequate speed of rotation such as 1000 rpm is used [94]: Either, an attacker would have to slow down the mesh so they can perform a manual attack, or they

would have to use a robot. The first class of attack would require the attacker to falsify the readings of the centrifugal accelerometer. MEMS accelerometers are complex devices, and the simplest way to falsify its readings would be to attach a circuit to the accelrometer's data bus that overrides the measurement result data. Creating such a circuit is easy, the challenge the attacker would have to overcome would be to access this bus and attach this circuit to the mesh in motion without stopping or disturbing it. At high speeds, this would necessarily require a custom attack robot.

4.2 Contactless attacks on the payload

Contactless attacks such as electromagnetic (EM) side-channel attacks or optical fault injection attacks on the payload could conceivably be conducted from the outside of the mesh. The efficacy of EM side-channel as well as fault injection attacks decays quickly with increased distance between probe and target, and they can be counteracted by simply placing the QKD relay's components such that they are spaced apart from the mesh. Optical attacks, on the other hand can be carried out even at a distance using appropriate focusing optics. The easiest way to prevent such attacks would be to place the payload into an opaque enclosure inside the mesh.

An additional variant of optical attacks would be using a laser to cut or drill into the payload. Such attacks can be impeded through several defense-in-depth measures. First, the payload QKD relay should be designed such that destroying any part of it such as connecting wires or fibers causes it to fail secure. Irrespective of attacks, this is a reasonable design objective anyway given that components could fail, and a component failure should never put the device in an insecure state. Further, similar to other optical attacks, a shield can be used to prevent laser cutting or drilling attacks as well with the only difference being the kind of shield. To prevent laser cutting or drilling, a thick metal shield can be used. The large thermal mass, high thermal conductivity and reflective surface of such a shield makes it difficult to cut. There are lasers such as pulsed Nd:YAG lasers that can cut even thick steel, but these this cutting produces a large amount of metal plasma and debris, which would likely destroy the payload in the process.

To make sure any active laser attack is quickly detected, as a final line of defense, both mesh and payload should include wideband optical sensors in their array of environmental tamper sensors. For instace, high-power pulsed lasers do not deposit much heat into their target because the surface

of the target is vaporized by the laser pulse too quickly, and thus might not trigger a simple temperature alarm inside the payload. In contrast, optical sensors even outside of the laser's wavelength range would have no trouble detecting the light emitted from the metal plasma created by the laser's pulses on impact with the payload.

4.3 Fast, mechanical attacks on the payload

A final class of attacks are mechanical attacks where an attacker mechanically compromises the IHSM QKD relay so quickly that the tamper alarm mechanism has no time to act. An instance of such an attack would be using a gun to fire a bullet at the payload, aiming to selectively destroy parts of it that are involved in tamper alarm response before they can act. This class of attack can be counteracted in similar ways as the previously mentioned optical attacks. Destruction of parts of the payload should never let it fall into an insecure state, meaning that such an attack alone should never be enough to compromise the QKD relay. There is little one can do to prevent destruction of the payload by projectile or by explosive, but a thick metal shield around the payload would make it more difficult to selectively target part of it using a projectile.

5 Outlook

5.1 Achievable security guarantees

Like conventional HSMs, Inertial HSMs are only ever an engeineering answer to a security question. In contrast with cryptographic solutions that can achieve provable, information-theoretic security in some cases, an IHSM's security rests upon an assumption on the engineering capabilities of an attacker. In contrast to conventional HSMs, which achieve this engineering assumption through the manufacture of hard-to-manipulate tamper sensing meshes, Inertial HSMs achieve it by rotating their tamper sensing mesh. In a conventional HSM, increasing the security of the tamper sensing mesh requires fine-tuning a bespoke manufacturing process. In contrast, increasing the security of an IHSMs simply requires making the rotor faster.

5.2 Trust bootstrapping

A key question in any trusted hardware deployment is how to bootstrap trust in a new device when faced with the possibility of supply-chain attacks. Conventional HSMs are only manufactured by a single manufacturer, and the common solution is to just trust that manufacturer. The HSM's manufacturer can factory-provision an identity key to the HSM that can be used to ascertain the HSM's integrity during shipping to the customer.

One of the key components of IHSM technology is that it does not require specialized components, or potting of the payload. While an IHSM could be manufactured and sold as a complete unit like a conventional HSM, their more modular nature makes it possible to place more control in the IHSM's customer. In particular, an IHSM could be sold without a payload installed, leaving the customer to install their own payload (such as a QKD node) inside the IHSM. Like a conventional HSM, the IHSM could be run during shipping to detect supply-chain attacks. Going further, since IHSMs are build from commodity components, the user could directly license the IHSM design and manufacturer it themselves, given them full control over the hardware supply chain. In a QKD deployment, the manufacturer of the QKD node could build both the QKD subsystem and the IHSM and integrate both, given that this would not require additional manufacturing capabilities due to the IHSM's simple construction.

5.3 Network implementation

IHSM-secured QKD nodes could be used to build QKD networks. IHSM-secured QKD nodes augment QKD network techniques such as Cao et al. [37], who present a network structure that exploits MDI-QKD to replace some of the network's nodes by untrusted nodes that do not require physical security.

5.4 Device Longevity

In any HSM application, failure of a single HSM must be mitigated through a backup and redundancy strategy that is carefully chosen such that it does not pose a security risk. Conventional HSMs are often operated in a cluster made from multiple HSMs. These clusters serve two purposes. First, they can compensate for the failure of a single HSM, which is crucial given that ideally, the HSM's secrets should never be stored outside the HSM. Second,

they improve processing rate by sharing load across their constituent HSMs. Since conventional HSMs are highly limited in their processing speed due to size and power dissipation constraints, this capacity is essential for some applications.

A cluster of Inertial HSMs can be set up in much the same way. In a QKD system, one implementation would be to run multiple QKD links in parallel. The secret key streams of all links could then be combined using a hash function like it is used in a single QKD link's privacy amplification step. When one QKD link fails, in this construction its secret key stream can safely be replaced by a stream of zeros as long as the remaining operating links in sum still provide sufficient entropy.

In an application where the overhead of multiple QKD links each requiring their own dark fiber would be too expensive, multiple IHSM-protected QKD transceivers could be connected to a single optical fiber through an optical switch. Mirco-Electromechanical Systems (MEMS)-based optical switchs are a well-established technology and can switch optical fibers within milliseconds at an insertion loss of no more than a decibel or two. In a QKD application, this insertion loss would be tolerable. Since QKD secret key rates stem from a stochastic process and as such are not constant, QKD systems buffer secret key bits. The switchover time of an optical switch used for failover between two QKD transceivers as well as the link establishment time of the failover transceiver can be absorbed by simply sizing this buffer appropriately.

6 Conclusion

In this chapter, we applied the Inertial Hardware Security Module (IHSM) concept to physically trusted relay nodes in a Quantum Key Distribution network. We note that the hardest challenge in the adoption of IHSMs in QKD relays is the fiber-optic passthrough between the outside world and the IHSMs QKD relay payload. We show three concepts along the spectrum trading off security and implementation complexity. All three concepts utilize a secondary rotating mesh on the inside of the primary mesh's shaft opening. We practically demonstrate one of our concepts, the offset labyrinth mesh, in a functional mechanical prototype. We experimentally measured the increase in loss of a standard telecommunications fiber when inserted through our mechanical prototype's fiber passthrough,

resulting in an increase in loss compared to a straight fiber that was below our measurement threshold of approximately 0.25 dB.

Web sources

[W50] *Corning SMF-28 Ultra Optical Fiber Product Information Sheet*. 2024-02. URL: https://www.corning.com/media/worldwide/coc/documents/Fiber/product-information-sheets/PI-1424-AEN.pdf (visited on 2024-09-05) (cit. on pp. 200, 202).

[W74] FS. *1M 12F SC/APC Singlemode Farbcodiertes LWL-Pigtail - FS.com Deutschland*. FS.com. URL: https://www.fs.com/de/products/42416.html (visited on 2024-09-05) (cit. on pp. 200, 202).

[W210] *Product Page: Fiber Pigtail LC/APC OS2 G652D-Yellow 2m - 1 Piece | Unique | O0485.2*. EFB-Elektronik GmbH. URL: https://www.efb-elektronik.de/en/fiber-pigtail-lc-apc-os2-g652d-yellow-2m-1-piece/o0485.2 (visited on 2024-09-05) (cit. on pp. 200, 202).

References

[11] Lyubov V. Amitonova et al. "Quantum Key Establishment via a Multimode Fiber". In: *Optics Express* 28.5 (2020-03-02), pp. 5965–5981. DOI: 10.1364/OE.380791 (cit. on p. 196).

[37] Yuan Cao et al. "Hybrid Trusted/Untrusted Relay-Based Quantum Key Distribution Over Optical Backbone Networks". In: *IEEE Journal on Selected Areas in Communications* 39.9 (2021-09), pp. 2701–2718. DOI: 10.1109/JSAC.2021.3064662 (cit. on pp. 197, 211).

[38] Yuan Cao et al. "The Evolution of Quantum Key Distribution Networks: On the Road to the Qinternet". In: *IEEE Communications Surveys & Tutorials* 24.2 (2022), pp. 839–894. DOI: 10.1109/COMST.2022.3144219 (cit. on p. 197).

[40] Wouter Castryck and Thomas Decru. "An Efficient Key Recovery Attack on SIDH". In: *Advances in Cryptology – EUROCRYPT 2023*. Ed. by Carmit Hazay and Martijn Stam. Vol. 14008. Cham: Springer Nature Switzerland, 2023, pp. 423–447. DOI: 10.1007/978-3-031-30589-4_15 (cit. on p. 193).

[44] José Chesnoy. *Undersea Fiber Communication Systems*. Second edition. Amsterdam: Academic Press, 2015. ISBN: 978-0-12-804269-4 (cit. on p. 196).

[78] Kathrin Garb et al. "FORTRESS: FORtified Tamper-Resistant Envelope with Embedded Security Sensor". In: *2021 18th International Conference on Privacy, Security and Trust (PST)*. 2021 18th International Conference on Privacy, Security and Trust (PST). 2021-12, pp. 1–12. DOI: 10.1109/PST52912.2021.9647783 (cit. on pp. 118, 198).

[79] Kathrin Garb et al. "The Wiretap Channel for Capacitive PUF-Based Security Enclosures". In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2022-06-08), pp. 165–191. DOI: 10.46586/tches.v2022.i3.165-191 (cit. on p. 198).

[80] Kathrin A Garb. "Tamper-Sensitive Design of PUF-Based Security Enclosures" (cit. on pp. 3, 115, 118, 124, 198).

[94] Jan Sebastian Götte and Björn Scheuermann. "Can't Touch This: Inertial HSMs Thwart Advanced Physical Attacks". In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2022), pp. 69–93. DOI: 10.46586/tches.v2022.i1.69-93 (cit. on pp. 75, 194, 200, 208).

[108] Wuzhou Hong et al. "Design and Compensation Control of a Flexible Instrument for Endoscopic Surgery". In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. 2020 IEEE International Conference on Robotics and Automation (ICRA). 2020-05, pp. 1860–1866. DOI: 10.1109/ICRA40945.2020.9196955 (cit. on p. 203).

[117] Vincent Immler et al. "Secure Physical Enclosures from Covers with Tamper-Resistance". In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2018-11-09), pp. 51–96. DOI: 10.46586/tches.v2019.i1.51-96 (cit. on pp. 37, 115, 141, 198).

[119] R. Impagliazzo. "A Personal View of Average-Case Complexity". In: *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*. Structure in Complexity Theory. Tenth Annual IEEE Conference. Minneapolis, MN, USA: IEEE Comput. Soc. Press, 1995, pp. 134–147. DOI: 10.1109/SCT.1995.514853 (cit. on p. 191).

[135] Joonhwan Kim et al. “Advancement of Flexible Robot Technologies for Endoluminal Surgeries”. In: *Proceedings of the IEEE* 110.7 (2022-07), pp. 909–931. DOI: 10.1109/JPROC.2022.3170109 (cit. on p. 203).

[150] Eufemia Lella and Giovanni Schmid. “On the Security of Quantum Key Distribution Networks”. In: *Cryptography* 7.4 (4 2023-12), p. 53. DOI: 10.3390/cryptography7040053 (cit. on p. 197).

[197] 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, 2018-06-24, pp. 1–6. DOI: 10.1145/3195970.3195976 (cit. on pp. 118, 124, 198).

[229] Ross T. Schermer and James H. Cole. “Improved Bend Loss Formula Verified for Optical Fiber by Simulation and Experiment”. In: *IEEE Journal of Quantum Electronics* 43.10 (2007-10), pp. 899–909. DOI: 10.1109/JQE.2007.903364 (cit. on p. 202).

[230] Andreas Schmitz et al. “A Rolling-Tip Flexible Instrument for Minimally Invasive Surgery”. In: *2019 International Conference on Robotics and Automation (ICRA)*. 2019 International Conference on Robotics and Automation (ICRA). 2019-05, pp. 379–385. DOI: 10.1109/ICRA.2019.8793480 (cit. on p. 203).

[243] Jung-wook Suh and Ki-young Kim. “Design of a Discrete Bending Joint Using Multiple Unit PREF Joints for Isotropic 2-DOF Motion”. In: *International Journal of Control, Automation and Systems* 15.1 (2017-02-01), pp. 64–72. DOI: 10.1007/s12555-016-0474-z (cit. on p. 203).

[269] Shuang Wang et al. “Twin-Field Quantum Key Distribution over 830-Km Fibre”. In: *Nature Photonics* 16.2 (2022-02), pp. 154–161. DOI: 10.1038/s41566-021-00928-2 (cit. on p. 196).

[280] Chao Yang, Hongqi Zhang, and Jinhai Su. “Quantum Key Distribution Network: Optimal Secret-Key-Aware Routing Method for Trust Relaying”. In: *China Communications* 15.2 (2018-02), pp. 33–45. DOI: 10.1109/CC.2018.8300270 (cit. on p. 197).

Chapter 8

Case Study: Multiparty Computation in Scalable Hardware Security Modules

> We can only desire based on what we know. It is our present experience of what we are and are not able to do that largely determines our sense for what is possible. This is why same sex relationships, in violation of sodomy laws, were a necessary precondition for the legalization of same sex marriage. This is also why those maintaining positions of power will always encourage the freedom to talk about ideas, but never to act.
>
> – *Moxie Marlinspike [W164], see also Rogaway [220]*

Contents

Inertial Hardware Security Modules do not only support much larger payloads compared to conventional HSMs, they also support much higher power dissipation since they allow for direct air cooling of their payload. Because they rotate at high speed, IHSM meshes do not need to be contiguous to provide adequate security. While a non-contiuous rotating mesh might theoretically allow a stationary attack tool to quickly penetrate, then retract through one of the mesh's gaps while the mesh is rotating, the time available for such an attack would be too short for a practical attack. For a mesh with three vertical connecting segments (cf. Figure 4.5 in Chapter 4) rotating at 1000 rpm, this time would be in the order of 20 ms. Conventional HSM monitoring circuits often require a similar amount of time to react to an attack [195].

Similar to how the increase in payload *size* unlocks new applications such as the Quantum Key Distribution relay use case we presented in Chapter 7, this increase in sustainable power dissipation by a factor of several hundred also unlocks a number of new applications. Especially applications that require large amounts of computing power benefit from IHSM technology, as their needs fundamentally cannot be met by conventional HSMs.

One such application that does not translate to conventional HSMs due to its need for large amounts of computing power is Multiparty Computation (MPC). MPC is a cryptographic construct that allows several networked parties to jointly perform a computation in such a way that the inputs to the computation remain private to the parties providing them, and no single party must be trusted for the computation to produce the correct result. Conceptually, MPC is similar to a secret sharing scheme that shares not just data, but computation between untrusted parties. The computation primitive MPC offers is a cryptographic answer to the question of how to bootstrap trust in a computing system.

To do
In this chapter, cite academic publications and patents on HSM cooling!

We can deconstruct the problem of trust in computing into two largely disjunct parts: Establishing trust in a computing system during its creation is one, and maintaining this trust throughout its life is the other. For the second part of this problem, maintaining trust in a system once trusted, we have an ample supply of good methods such as encryption, authentication, and formally proven protocols. In contrast, establishing trust in a computing system is largely intractable and despite a large corpus of academic research on approaches such as hardware trojan detection and physicaly unclonable functions, only two approaches find practical adoption: In one,

we build the system ourselves from the ground up, making sure to leave no part vulnerable to third-party compromise. In the other, we go to a store and physically buy a randomly-chosen computer using cash, assuming that while an attacker can target any particular system, they cannot target all systems simultaneously and we give them too little time to target the system we buy.

A limitation of both approaches is that in either case, while the party creating or acquiring the system can trust it, they cannot prove its trustworthiness to other parties. MPC solves this issue by allowing every party to contribute their trusted system to the protocol, cryptographically bootstrapping common trust in the computation and its output[1].

1 Fast MPC and Slow HSMs

MPC is a uniquely powerful cryptographic primitive, yet it has still not found widespread practical adoption. This is because MPC is extremely resource-intensive to run. MPC protocols exist on a continuum trading off between extreme memory and bandwidth requirements on one end and intense computational requirements on the other end. At a first glance, MPC and Hardware Security Modules look like they would complement each other well, but HSMs cannot keep up with the intense computational requirements posed by MPC.

Using P-256 curve ECC key generation as a benchmark, commercially available HSMs are quoted to perform between 3500 and 22000 cryptographic operations per second [145, 250, 259]. Meanwhile, an MPC protocol doing something as simple as a single AES encryption, corresponding to 7000 logic gates [270], requires tens of thousands of cryptographic operations when performed in MPC. As a result, applying conventional HSMs to MPC at any practical scale is infeasible by multiple orders of magnitude. Literature on MPC commonly uses server hardware as a platform for benchmarks.

HSMs are slow compared to contemporary computers because they are limited in their power dissipation, and power dissipation is largely proportional to processing speed. In the limited fields where HSMs have found

[1] In fact, MPC does more than just bootstrapping from each participant trusting their own system to a trusted shared computation. In an MPC protocol providing semi-honest or better security, MPC even *relaxes* each party's trust requirement from trusting their own system to trusting that any n-of-k out of all systems contributing to the protocol.

To do
Can we find a citation here?

commercial application, this limitation was never considered important and market forces pushing towards faster HSMs remain light . Fundamentally, conventional HSMs must envelope the entire payload in a tamper sensing mesh to detect drilling attacks, but a tamper sensing mesh that is impermeable to a drill is also impermeable to air. As a result, any heat conducted from the HSMs processor to the outside world must pass through the mesh. At the same time, the mesh cannot be thinned either because thinning it would enable micro-drilling attacks. The result of these constraints is a high thermal resistance between the HSM's processor and an external heat sink, which limits maximum power dissipation to a fraction of what is achieved in modern CPUs or even GPUs.

A secondary limitation of conventional HSMs is that the highly specialized tamper sensing foils used in their construction often cannot be scaled to arbitray sizes without incurring unsustainable process yields due to the multiplication of error rates with increasing area. As a result, even if the heat dissipation problem could be solved, manufacturing the tamper sensing foil for a conventional HSM large enough to contain a more powerful CPU might not be possible. The HSM's tamper-sensing envelope would have to protect not only the CPU itself, but also its supporting components such as memory, power supplies and any internal heat spreading components.

Inertial HSMs solve this issue since they allow their payload to be air cooled without compromising security, and they expand the feasible security boundary size from the several hundred milliliters offered by conventional HSMs to several liters and more, enabling the integration of standard, off-the-shelf server components such as mainboards, CPUs, CPU coolers, and power supplies. In this chapter, we will first provide a short overview of the theory of MPC before elaborating a design of an IHSM tailored to MPC tasks including performance calculations and unique design aspects. We will conclude with an outlook of applications unlocked by our design as well as promising areas for future improvements of our design.

2 The Fundamentals of Multiparty Computation

Secure Multiparty Computation can be separated into two broad classes of approaches: Garbled Circuits, and Secret Sharing-based techniques. Garbled Circuit techniques model the computation as a circuit of binary logic components such as logic gates. They are well-suited for implementing cryp-

tographic primitives such as conventional symmetric ciphers such as AES or hash functions such as the SHA-2 series. Secret Sharing-based techniques model computation as an arithmetic circuit made from components such as arithmetic operations. While they can also work in binary, they often support operations on larger finite fields. Secret sharing-based techniques are efficient processing integer numbers, but can have higher overhead in processing using many bitwise operations such as ciphers or cryptographic hash functions[66].

2.1 Security Models in MPC

MPC schemes are usually evaluated assuming one of three adversary levels: *Semi-Honest*, *Covert* or *Malicious* adversaries. A *Semi-Honest* adversary is an adversary that follows the protocol as specified, but that outside the protocol's execution may collude arbitrarily with other parties to reveal the secret inputs of other parties. A *Covert* adversary is an adversary that additionally may cheat during the protocol's execution, but only in ways that cannot be detected by other parties. Finally, a *Malicious* adversary is one that can deviate from the protocol's execution arbitrarily [17]. The covert adversary model most closely captures the requirements of a real-world scenario where a small number of cooperating parties runs the protocol, since in such settings cheating parties can easily be excluded once identified. The malicious adversary model captures real-world settings where parties do not have stable identities such as peer-to-peer settings. The semi-honest model is mostly interesting as a research tool since protocols assuming a semi-honest adversary can often be upgraded to covert or malicious security at some performance tradeoff. In a practical setting, a semi-honest secure MPC protocol would not provide additional security over just having one party run the computation except in some situations where inadvertent side-channel leakage is a concern.

2.2 Oblivious Transfer

Before we can go into details of multiparty computation, we need to define an important primitive. Oblivious Transfer is a cryptographic protocol between two parties that enables one party, the *Sender*, to share one of two values to the other party, the *Receiver*. The Receiver can choose which of the two values it wants to receive by inputting a choice bit b into the protocol. After the protocol has been executed, the Receiver learns *only*

its chosen value, and learns nothing about the other input value that the Sender provided. Meanwhile, the Sender learns nothing about the choice bit b of the receiver.

OT extensions

Oblivious Transfer is a public-key primitive, requiring asymmetric cryptographic operations. For this reason, "raw" Oblivious Transfer is not particularly fast. In practice, this issue is solved by applying a technique called Oblivious Transfer Extensions (OTe)[124]. Using OTe, a fixed, small number of public-key base Oblivious Transfer instances can be extended into an arbitrarily large number of Oblivious Transfer instances using only invocations of a pseudo-random function (PRF) such as a cryptographic hash function.

2.3 Boolean MPC with Yao's Garbled Circuits

Yao's Garbled Circuits (GC) protocol is one of the oldest Multiparty Computation protocols, dating back to the 1980ies. In Yao's GC, two parties jointly compute a function that is represented as a circuit of binary logic gates by evaluating the circuit gate by gate. In Yao's GC, one party, generator, creates a random *garbled* representation of the circuit and sends it to the other party, the evaluator, who computes its output. The core idea in Yao's GC is that every wire w_i in the circuit is assigned two random cryptographic secret keys w_i^b, called wire labels, one w_i^0 for the logical value 0 and one w_i^1 for the value 1. The mapping from logic values to these keys is assigned randomly by the generator, and unknown to the evaluator [281, 23, 66].

Gates are represented in Yao's GC as truth tables with one row for every combination of input wire values. Each row of these truth tables contains the output wire label (i.e. secret key) corresponding to the gate's logical output value for the row's combination of input values. The output wire label in each row is encrypted with *both* input wire labels corresponding to this row as keys.

The generator must indicate to the evaluator which row of a gate's truth table to decrypt, while also avoiding leaking the logical value of the output wire to the evaluator. This is commonly done in a technique called *point-and-permute* where a random pointer bit p_i^b is appended to each wire w_i^b label such that $p_i^b = \neg p_i^{\neg b}$. The rows in the gate's truth table are ordered

according to the combination of the two input wire labels' pointer bits. When the evaluator obtains the two input wire labels, they obtain their pointer bits, which combined are the index of the row to decrypt in the following gate's truth table.

It is clear how the protocol described above can be used to compute any binary circuit, but there are two questions remaining: How do the two parties provide input into the circuit, and how do they decode the circuit's output? Output is handled in Yao's GC by creating an output decoding table for every output wire of the circuit. The output decoding table contains two rows, one for a logical 0 output value and one for a logical 1 output value. Each row contains the hash of the output wire's label corresponding to the row's logical output value. This way, the evaluator can identify the logical value knowing the output wire label, but is unable to deduce the output wire label from the output decoding table.

Inputs are a bit more difficult to handle. While the generator can easily provide secret inputs by simply providing the evaluator with the input wire labels corresponding to its input, inputs from the evaluator require oblivious transfer to avoid leaking the evaluator's input to the generator. To input the logic bit b on wire w_i, the generator and the evaluator perform an 1-out-of-2 oblivious transfer with the generator assuming the Sender role and providing the two input wire labels w_i^0 and w_i^1 as the two choices, and the evaluator submitting its chosen input bit b as the OT's choice bit.

Yao's GC has good performance since the only asymmetric cryptographic primitive used is in the Oblivious Transfer needed for the evaluator's hidden inputs. The generation and the evaluation of the garbled circuit itself both only require evaluations of a pseudorandom function such as a cryptographic hash or a cipher. Still, performing a computation using a Garbled Circuit is many times slower than performing it in the clear. Intuitively, each single-bit gate in the garbled circuit results in several cryptographic operations with input and output sizes of dozens or hundreds of bits. Practically useful functions such as AES encryption have circuit implementations measuring thousands or tens of thousands of gates, meaning these costs quickly escalate for practical problem sizes [32, 238].

Count	Component	Power Dissipation (approx.)	Total
1	CPU	350 W [255]	350 W
16	Memory [W132]	2 W	32 W
1	Losses	40 W	40 W

Table 8.1: Power budget of a modern mid-range server. Losses were estimated at 10%, consistent with mainboard losses plus losses from a 80plus platinum efficiency certified power supply (94% at load).

3 A High-Performance IHSM for MPC Applications

Multiparty Computation is at the verge of being practical in some applications, but is still too computationally expensive for others. While some attempts at GPU-accelerating MPC primitives exist, in practice it is commonly implemented using CPU processing. The technology comes with an unavoidable increase in computational complexity since each single plaintext computation or gate results in several cryptographic operations.

A naive implementation might attempt to implement MPC using an HSM by simply offloading all cryptographic operations to the HSM. In practice, this is not a workable solution due to the slow processing speed of conventional HSMs. Conventional HSMs use low-power embedded processors since their encapsulation using potting and security meshes impedes heat transfer, limiting power dissipation.

In the near term, absent radical developments in either MPC theory or in the speed and power efficiency of processing hardware, the only feasible solution for HSM-protected MPC at any practical scale is to find a way to protect an entire server-class computer. As elaborated above, IHSMs are a natural fit for this requirement since they allow for large, air-cooled payloads.

To do
Refer to performance numbers from research above here

As a baseline performance target, we consider a commodity server mainboard in CEB or ATX form factor, populated with a high-end server CPU and a large amount of RAM. As MPC systems do not usually require a great amount of storage, we can largely ignore storage for our size and power calculations.

As a result, we end up with a total maximum power dissipation of approximately 420 W as shown in Table 8.1. Dissipating this amount of power using air cooling is within the capabilities of commodity server cooling components [51].

A common type of side-channel attack on cryptographic systems are power analysis attacks. In such attacks, the supply current of the target processing system is measured at high speed while the target is performing cryptographic computations. By aggregating the results of a large number of the resulting power traces, it is often possible to infer the value of secret data such as cryptographic keys. To mitigate this type of attack, not only do we have to place the CPU, mainboard, and memory inside of the HSM's tamper-sensing barrier, but also the power supply. A secondary benefit of placing the power supply inside the tamper-sensing barrier is that it simplifies the power wiring between the outside of the IHSM cage and the payload. Supplying the 12 V power rails that commodity mainboard commonly use requires tens of Ampere. To carrie such high current, the wiring has to be sized accordingly. In an IHSM, even thick wires can easily be passed through the mesh cage, but such wiring requires a large opening at the shaft on one end of the cage, which creates a literal security gap. Placing the power supply inside of the cage reduces the size of the wires needed since the power supply steps down a lower current 240 V input to the system's high-current 12 V rails. According to DIN VDE 0298-4 , a pair of 1.5 mm^2 conductors is sufficient for more than 3 kW of load under worst-case conditions.

To do
Citation?

3.1 Software Considerations

While the hardware of a HSM-assisted MPC system is a straightforward application of IHSM technology to a server platform, the software implementation poses some unique challenges. A core concern in an IHSM based on commodity hardware running a commodity operating system is the concrete implementation of the IHSM's alarm reseponse. When the IHSM detects tampering, it is crucial that all secrets in the payload have been made unusable before an attacker can either extract them, or stop the system from making them unusable.

Making secret data unusable to an attacker can be done by either deleting it (*zeroization* in HSM terminology) or by encrypting the data when it is stored and destroying the key. Zeroization is more technically challenging, while encryption comes at a performance cost. The main challenges in zeroization are ensuring that the data is deleted fast enough, and making sure it cannot be reconstructed through data remanence effects. Zeroization is practical for small amounts of RAM such as a microcontroller's main

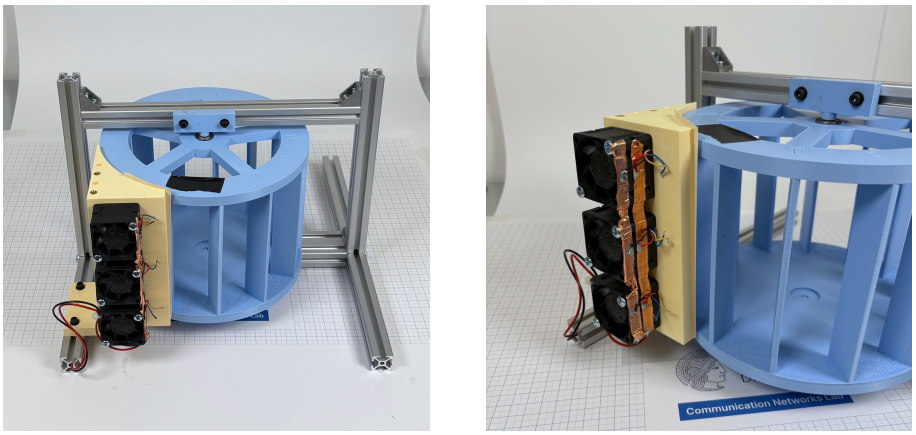

(a) (b)

Figure 8.1: Conceptual demonstrator of the fan-driven IHSM primary mesh approach.

SRAM or a small amount of DRAM in a server set aside for cryptographic keys. For larger amounts of data, or for any data stored on flash memory, HDDs etc., encryption is the only viable option due to speed. Encrypting data at rest on HDD or flash storage is straightforward to do in software and is a standard feature in any modern operating system. RAM is more challenging. Encrypting data in RAM cannot be done in software without effectively running a system emulation and incurring a massive performance penalty. Recent CPUs from Intel and AMD contain hardware features that provide transparent DRAM encryption. These hardware features would be necessary when securing an entire sever in an MPC setup with IHSMS technology.

3.2 A Joint Cooling and IHSM Envelope Powertrain

In this section, we will present a sketch of a design for an IHSM envelope large enough to fit a small server mainboard, and that provides air cooling to the payload. Our sketch solves the engineering issue of moving such an IHSM envelope while simultaneously providing cooling to the payload.

Our proposed design is based on the idea of using the cooling fans' airflow to power the rotation of the IHSM envelope. Figure 8.1 shows a conceptual demonstration of this concept. Using the basic cylindrical design, the IHSM envelope consists of two discs above and below the payload that are connected through vertical struts containing part of the tamper-sensing mesh on the outside of the payload. We propose widening these vertical connecting struts, and angling them such that the entire envelope

becomes a centrifugal impeller. By letting air flow into the envelope from the side, and back out through its top and bottom, the envelope assumes the same configuration used in centrifugal cooling fans.

Laying out an IHSM this way has several advantages. First, we save some vertical space by removing the motor from the shaft of the mesh. Second, on top of driving the mesh, the airflow also serves to cool the payload. Second, this approach eliminates the motor driving the mesh as a single point of failure. In a basic IHSM design as we introduced it in Chapter 4, this motor is a critical component as it failing would lead to the mesh accelerometer triggering the tamper alarm. Using a brushless motor type the numebr of wear components in this motor can be reduced to the motor's shaft bearings. A complication in the practical manufacturing of IHSMs at a small scale is that small-scale production does not allow for a custom-made motor. Limiting the selection to off-the-shelf brushless motors leads to an unpredictability of bearing life since precise bearing specifications are not usually included in motor datasheets.

Compared to the market for off-the-shelf small brushless motors, cooling fans are easier to shop for. A large selection of products with various form factors and specifications is available, and manufacturers usually give detailed information on both performance and lifetime. For industrial and server cooling fans, uninterrupted 24/7 operation is their nominal operating condition.

The cooling fans can be located on the outside of the envelope in an easily accessible location, and can be set up in a redundant way such that a failed cooling fan can be replaced while the system continues operation.

The main drawback of a fan-driven IHSM is the necessary airflow. To maximize payload volume, the fan blades must be kept as narrow as possible. Narrow fan blades work best at high air speed, but high air speed requires the fan to have high airflow. Besides limiting fan selection and increasing power consumption, high airflow fans also are noisy. Despite these limitations, we consider fan-driven IHSMs a valid tradeoff since such a system would most likely be deployed in a datacenter where high noise levels are acceptable.

4 Outlook

In this chapter we briefly introduced the challenges raised by MPC at scale, and we outlined a practical solution based on an IHSM-protected sever that can be used to perform MPC with an unique combination of high bandwidth and low latency that was previously considered impractical due to physical security concerns.

Web sources

[W132] Patrick Kennedy. *DDR4 DIMMs and System Power Consumption - We Tested*. ServeTheHome. 2017-01-30. URL: https://www.servethehome.com/ddr4-dimms-system-power-consumption-tested/ (visited on 2025-10-27) (cit. on p. 224).

[W164] Moxie Marlinspike. *We Should All Have Something To Hide*. Blog of Moxie Marlinspike. 2013-06-12. URL: https://moxie.org/2013/06/12/we-should-all-have-something-to-hide.html (visited on 2025-11-18) (cit. on pp. 2, 217).

References

[17] Yonatan Aumann and Yehuda Lindell. "Security Against Covert Adversaries: Efficient Protocols for Realistic Adversaries". In: *Journal of Cryptology* 23.2 (2010-04), pp. 281–343. DOI: 10.1007/s00145-009-9040-7 (cit. on pp. 2, 221).

[23] D. Beaver, S. Micali, and P. Rogaway. "The Round Complexity of Secure Protocols". In: *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing - STOC '90*. The Twenty-Second Annual ACM Symposium. Baltimore, Maryland, United States: ACM Press, 1990, pp. 503–513. DOI: 10.1145/100216.100287 (cit. on p. 222).

[32] Joan Boyar and René Peralta. "A New Combinational Logic Minimization Technique with Applications to Cryptology". In: *Experimental Algorithms*. Ed. by Paola Festa. Berlin, Heidelberg: Springer, 2010, pp. 178–189. DOI: 10.1007/978-3-642-13193-6_16 (cit. on p. 223).

[51] Vlad C Coroamă et al. *Past and Possible Future Trends*. International Energy Agency, 2025-04 (cit. on p. 224).

[66] David Evans, Vladimir Kolesnikov, and Mike Rosulek. "A Pragmatic Introduction to Secure Multi-Party Computation". In: () (cit. on pp. 221, 222).

[124] Yuval Ishai et al. "Extending Oblivious Transfers Efficiently". In: *Advances in Cryptology - CRYPTO 2003*. Ed. by Dan Boneh. Red. by Gerhard Goos, Juris Hartmanis, and Jan Van Leeuwen. Vol. 2729. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 145–161. DOI: 10.1007/978-3-540-45146-4_9 (cit. on p. 222).

[145] Dinesh Kumar. *IBM Z16 Performance of Cryptographic Operations: Cryptographic Hardware: CPACF, CEX8S with Quantum-Safe CRYSTALS Algorithms*. IBM, 2025-03 (cit. on p. 219).

[195] Johannes Obermaier and Vincent Immler. "The Past, Present, and Future of Physical Security Enclosures: From Battery-Backed Monitoring to PUF-based Inherent Security and Beyond". In: *Journal of Hardware and Systems Security* 2 (2018), pp. 289–296. DOI: 10.1007/s41635-018-0045-2 (cit. on pp. 2, 37, 63, 78, 218).

[220] Phillip Rogaway. "The Moral Character of Cryptographic Work". In: *Advances in Cryptology*. ASIACRYPT 2015. Vol. 9452 & 9453. LNCS. Auckland, New Zealand: Springer, 2015, p. XVIII. DOI: 10.1007/978-3-662-48800-3 (cit. on pp. 2, 22, 217).

[238] Ebrahim M. Songhori et al. "TinyGarble: Highly Compressed and Scalable Sequential Garbled Circuits". In: *2015 IEEE Symposium on Security and Privacy*. 2015 IEEE Symposium on Security and Privacy. 2015-05, pp. 411–428. DOI: 10.1109/SP.2015.32 (cit. on p. 223).

[250] *Thales Luna Network HSM Product Brief*. Thales, 2024-10 (cit. on p. 219).

[255] Hannes Tröpgen et al. "16 Years of SPEC Power: An Analysis of X86 Energy Efficiency Trends". In: *2024 IEEE International Conference on Cluster Computing Workshops (CLUSTER Workshops)*. 2024 IEEE International Conference on Cluster Computing Workshops (CLUSTER Workshops). 2024-09, pp. 76–80. DOI: 10.1109/CLUSTERWorkshops61563.2024.00020 (cit. on p. 224).

[259] *U.Trust General Purpose HSM Se-Series Datasheet*. utimaco, 2025-04 (cit. on p. 219).

[270] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. "Global-Scale Secure Multiparty Computation". In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. CCS '17: 2017 ACM SIGSAC Conference on Computer and Communications Security. Dallas Texas USA: ACM, 2017-10-30, pp. 39–56. DOI: 10.1145/3133956.3133979 (cit. on p. 219).

[281] Andrew Chi-Chih Yao. "How to Generate and Exchange Secrets". In: *27th Annual Symposium on Foundations of Computer Science (Sfcs 1986)*. 27th Annual Symposium on Foundations of Computer Science (Sfcs 1986). 1986-10, pp. 162–167. DOI: 10.1109/SFCS.1986.25 (cit. on p. 222).

Chapter 9

Conclusion

In this thesis, we propose Inertial Hardware Security Modules (IHSMs), a new approach to physical security that combines conventional tamper-sensing meshes with physical movement to bootstrap a highly secure system from low-security, off-the-shelf parts, solving our first research question introduced in Chapter 1. To motivate our research, we use the German national digital health record system as an example demonstrating the difficulties in achieving useful hardware security in practice. Besides some minor cryptographic oddities, our analysis reveals at least one essential specification mistake that negates the hardware security of the system by unnecessarily introducing a poorly protected HSM. With this motivation in mind, we support the construction of concretely secure IHSMs by providing deep analyses of two key engineering challenges in IHSM construction, mesh monitoring and power transfer. Solving our second research question, we propose a low-cost TDR-based mesh monitoring system that exceeds the capabilities of previous systems from academic or from patent literature. Our system is capable of monitoring large meshes while simultaneously providing detailed results. Our TDR-based mesh monitoring system is of independent interest, since it can also be integrated into traditional HSM designs. We additionally propose a new, generalized design for high-frequency PCB inductors with low parasitic capacitance. Our design provides better bandwidth and lower parasitic capacitance compared to the state of the art without increasing implementation cost. We conclude this thesis with two chapters elaborating on two new use cases that are made possible by IHSM technology due to its ability to protect large payloads that have high power consumption. Together, these results answer our third and final research question.

The research presented in this thesis is aimed at advancing both academic research and applied engineering in hardware security. We believe that by publishing our research including its artifacts under open source licenses, we provide the basis for future research in tamper-sensing technology, a field that remains under-served in today's academic landscape.

Recent history has shown that state-level adversaries are a mounting threat to civil rights organizations, human rights lawyers, members of minorities, and many others. While western democracies used to be considered safe havens of human rights, today human rights are under attack both from within and from the outside in countries across the globe. Publishing IHSM technology as open source, we hope to provide one building block for new computing systems accessible to all that are resilient and secure in the face of growing adversity.

Outlook

With the research contributions we presented in this thesis, we open up a new field of hardware security research centered on Inertial HSMs and improvements to conventional tamper sensing meshes. Below, we will list some research directions that we consider worthwhile for future investigation.

- Improving the resolution of the sampling mesh monitoring approach we presented in Chapter 5. Possible improvements include increasing pulse risetime through a discrete transistor amplifier circuit, as well as evaluating an FPGA as a replacement for the microcontroller to take advantage of the improved delay primitives offered by many FPGA families.

- Characterizing the PUF-like effects we observed in Chapter 5 in mesh coupons using our sampling mesh monitoring approach.

- Integrating IHSM technology with a HSM firmware implementation into a small form factor to create a portable IHSM. A small form factor introduces new challenges besides the mere integration of the necessary circuitry and placement of the mesh. For instance, wireless power and data transfer would need to be integrated with the device without disrupting mesh monitoring. An on-axis solution would likely require magnetic shielding materials and possible non-magnetic ceramic bearings. Furthermore, integrating a sufficiently small motor

and optimizing the design for long bearing life is challenging at the high rotation speed necessary at a small overall diameter. Finally, at high speeds, precisely balancing the whole assembly to avoid vibrations that could lead to early mechanical failure is difficult.

- Exploring IHSM applications beyond what we outlined in this thesis. For instance, one application of recent interests would be physically securing GPUs used for AI training. The background for such work could be either export control motivations, or a concern for security and privacy of user input, training data, or even trained weights.

We will proceed with future research into IHSM applications. We have published our results up to this point as open source hardware and software, and we intend to build on these publications.