IACR Transactions on Cryptographic Hardware and Embedded Systems
ISSN 2569-2925, Vol. 2022, No. 1, pp. 69–93. DOI:10.46586/tches.v2022.i1.69-93

Can't Touch This: Inertial HSMs Thwart Advanced Physical Attacks

Jan Sebastian Götte and Björn Scheuermann

HIIG, Berlin, Germany
ihsm@jaseg.de,bjoern.scheuermann@hiig.de

Abstract. In this paper, we introduce a novel countermeasure against physical attacks: Inertial Hardware Security Modules (IHSMs). Conventional systems have in common that their security requires the crafting of fine sensor structures that respond to minute manipulations of the monitored security boundary or volume. Our approach is novel in that we reduce the sensitivity requirement of security meshes and other sensors and increase the complexity of any manipulations by rotating the security mesh or sensor at high speed—thereby presenting a moving target to an attacker. Attempts to stop the rotation are easily monitored with commercial MEMS accelerometers and gyroscopes. Our approach leads to an HSM that can easily be built from off-the-shelf parts by any university electronics lab, yet offers a level of security that is comparable to commercial HSMs. We have built a proof-of-concept hardware prototype that demonstrates solutions to the concept's main engineering challenges. As part of this proof-of-concept, we have found that a system using a coarse security mesh made from commercial printed circuit boards and an automotive high-g-force accelerometer already provides a useful level of security.

Keywords: hardware security · implementation · smart cards · electronic commerce

1 Introduction

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 [New, Fra, JRR+18]. 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 [AHT+20, And].

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

Received: 2021-07-15 Accepted: 2021-09-15 Published: 2021-11-19

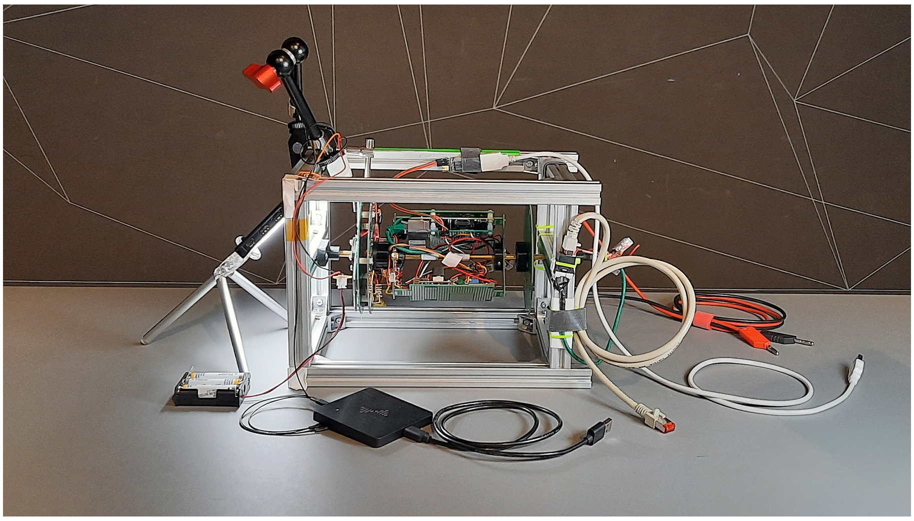

Figure 1: The prototype as we used it to test power transfer and bidirectional communication between stator and rotor. This picture shows the proof-of-concept prototype's configuration that we used for accelerometer characterization (Section 6) without the vertical security mesh struts that connect the circular top and bottom outer meshes.

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 [DMA08]. 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 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 paper 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 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 paper with a general evaluation of our design in Section 7.

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 [Boa]. 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 [IMFC13, ION$^+$, And]. There has been some research on monitoring the HSM's interior using e.g. electromagnetic radiation [TZP, KA12] or ultrasound [Vri] but none of this research has found widespread adoption yet.

HSMs can be compared to physical seals [And]. 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 Physically 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 optical 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) [Int].

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 [Boa], 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 [And], 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 [SW]. 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 [SW], 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 [OI18, DMA08, And, IMFC13].

Shifting our focus from industry use to the academic state of the art, in [ION$^+$], 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 monitoring. 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 [TZP], 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. [TZP] 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 [ION$^+$], at the cost of worse and less predictable security guarantees. Where [TZP] use electromagnetic radiation, Vrijaldenhoven in [Vri] uses ultrasound waves traveling on a surface acoustic wave (SAW) device to a similar end.

While Tobisch et al. [TZP] approach the sensing frontend cost as their primary optimization target, the prior work of Kreft and Adi [KA12] 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 [ION$^+$] while commercial vendors concentrate on means to certify and cheaply manufacture these security barriers [DMA08]. 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 [Rah] 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 [Hai] and is routinely used in military applications to make things harder to hit [Ter13] 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 paper, 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 Physically Unclonable Function as a means of tamper detection [TZP, ION$^+$]. 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 [KVK, SH, Cam, Eln].

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 $1000\,\mathrm{m\,s^{-1}}$ 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 2. Visible are the axis of rotation, an accelerometer on the rotating part that is used to detect 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

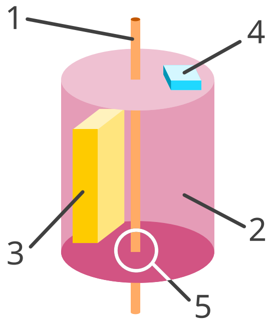

Figure 2: Concept of a simple spinning Inertial HSM. 1 - Shaft. 2 - Security mesh. 3 - Payload. 4 - Accelerometer. 5 - Shaft penetrating security mesh.

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 [IMFC13]. 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 [Obe]. 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 $0.1\,\mathrm{K\,W^{-1}}$ [Fyl]. If one were to make an HSM's security mesh out of an average thermally conductive epoxy with thermal conductivity $k \approx 1\,\mathrm{W\,m^{-1}\,K}$ [Kor98, Sha, MG], the resulting thermal resistance for a 5-by-5 centimeter, 5 mm thermal interface alone would be $2\,\mathrm{K\,W^{-1}}$, 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" [Gro], 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 [Hol]. Traditionally, this problem is solved by storing all secrets in multiple, geographically redundant HSMs [NV]. 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 [Lun] created by signal.org, the organization running the signal secure messenger app. In this demonstration, signal.org have implemented the Raft consensus algorithm [OO14] 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 [Obe]. 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 [Adm] 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 $\mathrm{km\,h^{-1}}$ and above [ika02, GCKMT]. Measurements of the Peak Ground Acceleration (PGA) of severe earthquakes show that even the strongest earthquakes rarely reach a PGA of 0.1 g [FT] 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 [Kel, Bea, Dix], 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 [Roo], 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 [Obe], 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 $\frac{5\,\mathrm{mm}\cdot 5\,\mathrm{mm}}{100\,\mathrm{mm}\cdot 100\,\mathrm{mm}} = 0.25\%$. 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 [Obe]. This design is vulnerable to attempts to stick a fine needle through the interface between lid and PCB [NBd]. 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 [NBd]. 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 [And]. 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 [And].

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

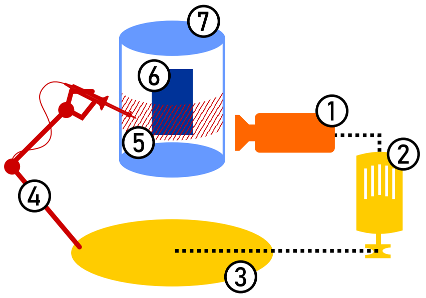

Figure 3: Schematic overview of a robotic rotating-stage attack. An optical sensor (1) observes the IHSM's rotation and adjusts the setpoint of a servo motor (2) that rotates the attack stage (3). On the rotating attack stage, a remote-controlled manipulator (4) is mounted that deactivates the security 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).

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 [BC19]. 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.

Figure 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 manipulation of the IHSM is hard. Let us assume a small IHSM

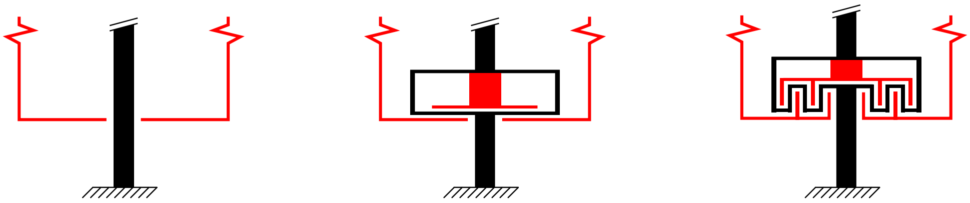

(a) Cross-sectional view of the basic configuration with no special protection of the shaft. Red: moving mesh – Black: stationary part.

(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: Mechanical countermeasures to attacks through or close to the shaft of a fixed-axis rotating 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 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 [SW]. 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 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 [NBd]. 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 [TWX+17], 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 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 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 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 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

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

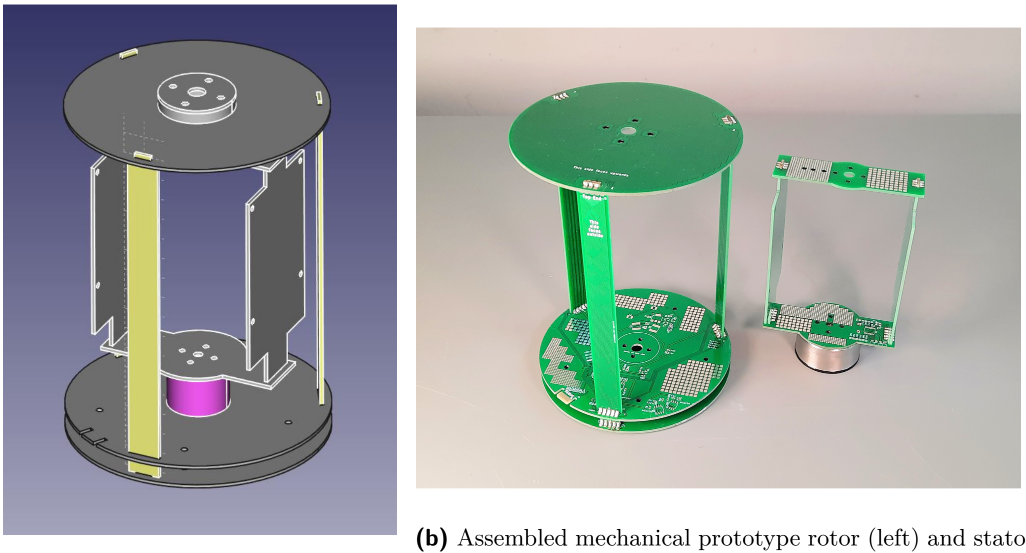

(a) The 3D CAD design of the proof-of-concept prototype.

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

Figure 5: Our proof-of-concept prototype IHSM's PCB security mesh design

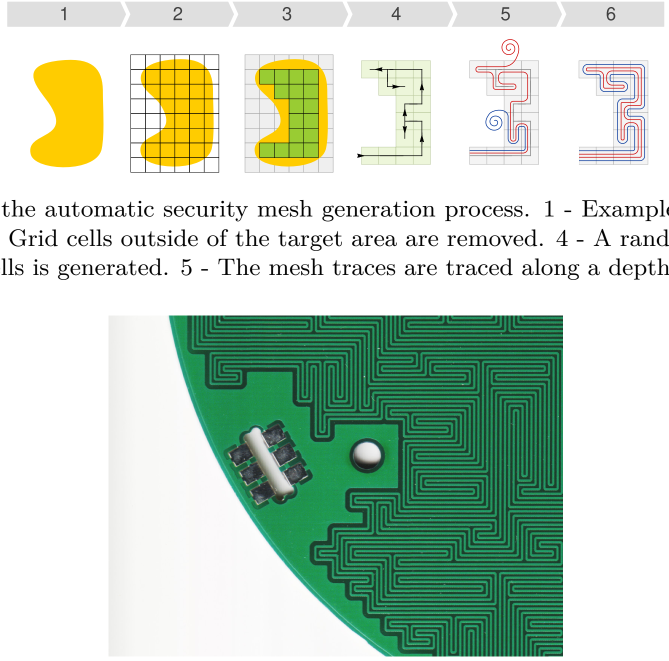

(a) Overview of the automatic security mesh generation process. 1 - Example target area. 2 - Grid overlay. 3 - Grid cells outside of the target area are removed. 4 - A random tree covering the remaining cells is generated. 5 - The mesh traces are traced along a depth-first walk of the tree. 6 - Result.

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

Figure 6: Our automatic security mesh generation process

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 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 sub-family 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 [Per]. While the initial handshake for key establishment uses elliptic-curve cryptography and may take several hundred milliseconds [TPGV], 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 $2\,\mu\text{A}\,\text{MHz}^{-1}$ 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 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 $\frac{1}{2}V_{\text{CC}}$ feeding into an `MCP6494` general purpose opamp configured as a 100 kΩ transimpedance amplifier. As shown in Figure 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 7a shows the

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

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

[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.

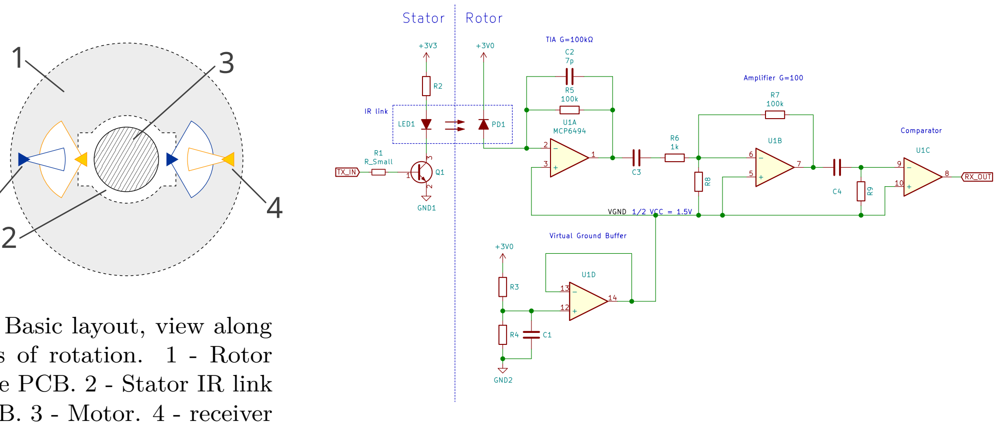

(a) Basic layout, view along axis of rotation. 1 - Rotor base PCB. 2 - Stator IR link PCB. 3 - Motor. 4 - receiver PIN photodiode. 5 - transmitter IR LED.

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

Figure 7: IR data link implementation

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.

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 ±120 g. At its 14-bit resolution, one LSB corresponds to 15 mg.

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 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 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 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.

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

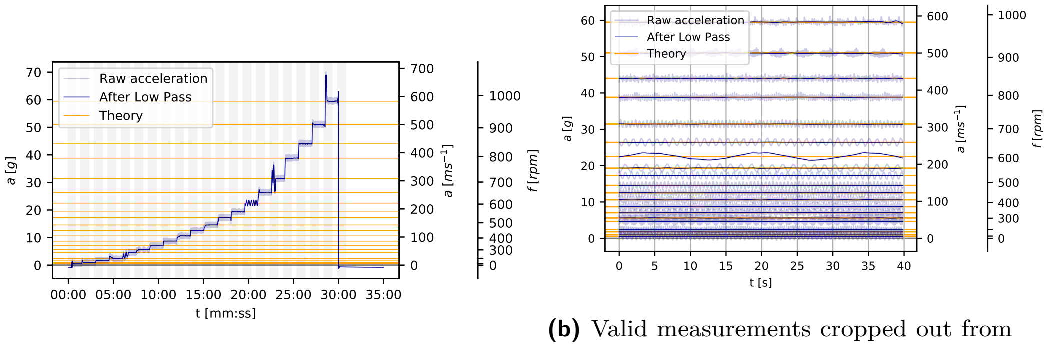

(a) Raw recording of accelerometer measurements during one experiment run. Shaded areas indicate time intervals when we manually adjusted speed.

(b) Valid measurements cropped out from 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 8: Traces of acceleration measurements during one experiment run.

Figure 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 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 paper, 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 security 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. We published all design artifacts of our PoC online, please refer to Appendix A for details. The next steps towards a practical application of our design

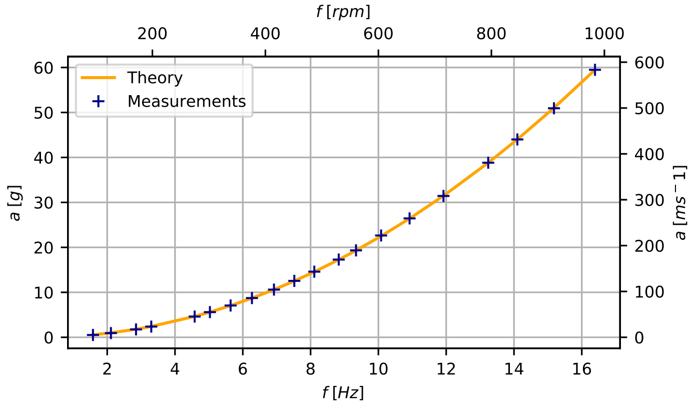

Figure 9: Centrifugal acceleration versus angular frequency in theory and in our experiments. Experimental measurements are shown after correction for offset and scale error. Above 300 rpm, the relative error is below 0.5 %. Below 300 rpm, the residual offset error has a large impact (0.05 g absolute or 8% relative at 95 rpm.)

will be to design a manufacturable stator/rotor interface with inductive power and data transfer integrated into the motor’s magnetics and a custom motor driver tuned for the application that is able to precisely measure both angular velocity and winding current for an added degree of tamper detection through the measurement of external forces acting on the rotor.

References

[Adm] US Federal Aviation Administration. Pack safe: Batteries, lithium.

[AHT+20] Nils Albartus, Max Hoffmann, Sebastian Temme, Leonid Azriel, and Christof Paar. DANA universal dataflow analysis for gate-level netlist reverse engineering. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2020(4):309–336, 2020.

[And] Ross Anderson. *Security Engineering*.

[BC19] 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.

[Bea] C. F. Beards. *Structural Vibration: Analysis and Damping*. Wiley.

[Boa] David G. Boak. A history of u.s. communications security, volumes i and ii. Lecture Notes.

[Cam] Bertrand Campagnie. Choose the right accelerometer for predictive maintenance. Technical report, Analog Devices.

[Dix] John C. Dixon. *The Shock Absorber Handbook*. Wiley.

[DMA08] 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)*, page 281–295. IEEE, 2008.

[Eln] Maged Elsaid Elnady. *On-Shaft Vibration Measurement Using a MEMS Accelerometer for Faults Diagnosis in Rotating Machines.* PhD thesis.

[Fra] Jessie Frazelle. Securing the boot process: The hardware root of trust.

[FT] Yoshimitsu Fukushima and Teiji Tanaka. A new attenuation relation for peak horizontal acceleration of strong earthquake ground motion in japan. 80:757 – 783.

[Fyl] Emmanouil D. Fylladitakis. Top tier cpu air coolers q3 2015: 9-way roundup review.

[GCKMT] A. German, J-L. Comeau, M.J. Shkrum K.J. McClafferty, and P.F. Tiessen. Event data recorders in the analysis of frontal impacts. In *Annual Proceedings of the Association for the Advancement of Automotive Medicine*, pages 225–243.

[Gro] Thales Group. Thales luna hsm product family overview page.

[Hai] Lester Haines. Us outfit patents 'invisible' uav: Stealth through persistence of vision.

[Hol] Martin Holland. Cloud-dienstleister ovh: Feuer zerstört rechenzentrum, ein weiteres beschädigt.

[ika02] A test procedure for airbags, 2002.

[IMFC13] Phil Isaacs, Thomas Morris Jr, Michael J Fisher, and Keith Cuthbert. Tamper proof, tamper evident encryption technology. Technical report, Surface Mount Technology Association, 2013.

[Int] International Atomic Energy Agency. *Safeguards, techniques and equipment*, volume 1 of *International Nuclear Verification Series.*

[ION$^+$] Vincent Immler, Johannes Obermaier, Kuan Kuan Ng, Fei Xiang Ke, Jin Yu Lee, Yak Peng Lim, Wei Koon Oh, Keng Hoong Wee, and Georg Sigl. Secure physical enclosures from covers with tamper-resistance. *IACR transactions on cryptographic hardware and embedded systems.*

[JRR$^+$18] Scott Johnson, Dominic Rizzo, Parthasarathy Ranganathan, Jon McCune, and Richard Ho. Titan: enabling a transparent silicon root of trust for cloud. In *Hot Chips: A Symposium on High Performance Chips*, 2018.

[KA12] Heinz Kreft and Wael Adi. Cocoon-puf, a novel mechatronic secure element technology. 2012.

[Kel] S. Graham Kelly. *Fundamentals of Mechanical Vibrations.* McGraw-Hill Series in Mechanical Engineering. McGraw-Hill, 2 edition.

[Kor98] Tony Kordyban. *Hot Air Rises and Heat Sinks: Everything You Know about Cooling Electronics is Wrong.* ASME, 1998.

[KVK] Ivar Koene, Raine Viitala, and Petri Kuosmanen. Internet of things based monitoring of large rotor vibration with a microelectromechanical systems accelerometer.

[Lun] Joshua Lund. Technology preview for secure value recovery.

[MG] MG Chemicals. Mg chemicals specialty adhesives catalog.

[NBd] Karsten Nohl, Fabian Bräunlein, and dexter. Shopshifting: The potential for payment system abuse.

[New] Lily Hay Newman. Apple's t2 security chip has an unfixable flaw.

[NV] Gemalto NV. Safenet pci-e hsm 6.2 product documentation: High availability (ha) overview.

[Obe] Johannes Obermaier. Physical unclonable functions: The future technology for physical security enclosures?

[OI18] 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. 2:289–296, 2018.

[OO14] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, pages 305–319, Philadelphia, PA, June 2014. USENIX Association.

[Per] Trevor Perrin. The noise protocol framework.

[Rah] Mujib Rahman. Optical fiber cable with tampering detecting means.

[Roo] Root Zone KSK Operator Policy Management Authority. Root zone ksk operator key management procedure.

[SH] 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. 16.

[Sha] Younes Shabany. *Heat Transfer: Thermal Management of Electronics*. CRC Press.

[SW] Sean Smith and Steve Weingart. Building a high-performance, programmable secure coprocessor. 31.

[Ter13] Daniel Terdiman. Aboard america's doomsday command and control plane, July 2013.

[TPGV] 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*.

[TWX$^+$17] Timothy Trippel, Ofir Weisse, Wenyuan Xu, Peter Honeyman, and Kevin Fu. Walnut: Waging doubt on the integrity of mems accelerometers with acoustic injection attacks. In *2017 IEEE European symposium on security and privacy*, page 3–18. IEEE, 2017.

[TZP] Johannes Tobisch, Christian Zenger, and Christof Paar. Electromagnetic enclosure puf for tamper proofing commodity hardware and other applications.

[Vri] Serge Vrijaldenhoven. Acoustical physical uncloneable functions.

A Source code and design artifacts

During our research on this paper, we have created a number of digital design artifacts including a 3D mechanical CAD model of our prototype IHSM, schematics, and PCB layouts for all of its PCBs including the prototype security mesh monitor PCB as well as firmware and data analysis scripts for the experiments we ran on the prototype IHSM. All of these digital artifacts as well as the sources to this paper are included in the git repository linked below. A copy of these design artifacts as well as our raw measurement data is included in the supplementary material to this paper.

`https://git.jaseg.de/ihsm.git`

This is version `v4.0-0-ga4184bb` of this paper, generated on November 17, 2021.