

Inertial HSMs Thwart Advanced Physical Attacks

Jan Sebastian Götte¹ and Björn Scheuermann²

¹ HIIG

ihsm@jaseg.de

² HU Berlin

scheuermann@informatik.hu-berlin.de

Abstract. In this paper, we introduce a novel countermeasure against physical attacks: Inertial hardware security modules (iHSMs). Conventional systems have in common that they try to detect attacks by crafting sensors responding to increasingly 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 a 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. By building prototype hardware we have demonstrated solutions to the concept’s engineering challenges.

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 has barely changed. Given the right skills, physical access to a computer still often means 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 [14, 6, 11].

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 [1, 2].

HSMs rely on a fragile foil with much larger-scale conductive traces being 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




Figure 1: The prototype as we used it to test power transfer and bidirectional communication between stator and rotor. In the picture, the prototype is missing the vertical security mesh struts connecting the circular top and bottom outer meshes that rotate around the stationary payload in the center.

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 [4]. Now consider the same HSM mounted on a large flywheel. In addition to its usual defenses the 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—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 3.1). 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 work 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 boundary sensing modes 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 on 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 the physical security of HSMs. 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 prototype system that we will illustrate in Section 5. 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 tracing back decades in their electronic realization. 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 [10, 8, 2]. There has been some research on monitoring the HSM’s inside using e.g. electromagnetic radiation [20, 13] or ultrasound [22] but none of this research has found widespread adoption yet.

HSMs can be compared to physical seals [2]. Both are tamper evident devices. The difference is that a HSM continuously monitors itself whereas a physical seal only serves to record tampering and requires someone to examine it. This examination can be by eye in the field, but it can also be using complex equipment in a laboratory. An HSM in principle has to have this examination equipment built-in.

Physical seals are used in a wide variety of applications, but the most interesting ones from a research point of view that are recorded in public literature are those used in monitoring of nuclear material under the International Atomic Energy Authority (IAEA). Most of these seals use the same approach that is used in Physically Uncloneable Functions, 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 seals lifetime, the device is returned from the field to the 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) [9].

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.

In [2], Anderson gives a comprehensive overview on physical security. An example they cite is the IBM 4758 HSM whose details are laid out in depth in [18]. 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 traditional construction of a flexible mesh wrapped around the module’s core. In [18], the authors state the module monitors this mesh for short circuits, open circuits and conductivity. The fundamental approach to tamper detection and construction is similar to other commercial offerings [16, 4, 2, 10].

In [8], Immler et al. describe a HSM based on precise capacitance measurements of a mesh. In contrast to traditional meshes, the mesh they use consists of a large number of individual traces (more than 30 in their example). Their concept promises a very high degree of protection. The main disadvantages of their concept are a limitation in covered area 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 [20], Tobisch et al. describe a construction technique for a hardware security module that is based around commodity Wifi hardware inside a conductive enclosure. In their design, an RF transmitter transmits a reference signal 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 RF cavity w.r.t. phase and frequency response. Their fundamental assumption is that the RF behavior of the cavity is inscrutable from the outside, and that even a small disturbance anywhere within the volume of the cavity will cause a significant change in its RF response. The core idea in [20] is to use commodity Wifi hardware to reduce the cost of the HSM's sensing circuitry. The resulting system is likely both much cheaper and capable of protecting a much larger security envelope than e.g. the design from [8], at the cost of worse and less predictable security guarantees. Where [20] use electromagnetic radiation, Vrijaldenhoven in [22] uses ultrasound waves travelling on a surface acoustic wave (SAW) device to a similar end.

While [20] approach the sensing frontend cost as their only optimization target, the prior work of Kreft and Adi [13] 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 HSM 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 [8] while commercial vendors concentrate on means to certify and cheaply manufacture these security barriers [4]. 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 an 1988 patent [17] that describes a mechanism to detect tampering along a communication cable by enclosing the cable inside a conduit filled with pressurized gas.

2.1 Patents literature

During development, we performed several hours of research on prior art for the inertial HSM concept. Yet, we could not find any mentions of similar concepts either in academic literature or in patents. Thus, we are likely the inventors of this idea and we are fairly sure it is not covered by any patents or other restrictions at this point in time.

Since the concept is primarily attractive for small-scale production and since cheaper mass-production alternatives are already commercially available, we have decided against applying for a patent and we wish to make it available to the general public without any restrictions on its use. We invite you to use it as you wish and to base your own work on our publications without any fees or commercial restrictions. Where possible, we ask you to cite this paper and attribute the inertial HSM concept to its authors.

3 Inertial HSM construction and operation

Mechanical motion has been proposed as a means of making things harder to see with the human eye [7] and is routinely used in military applications to make things harder to hit [19] 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: Rotation, pendulum, linear.

2. How to construct the **tamper detection mesh**.
3. How to **detect braking** of the HSM's movement.
4. The **mechanical layout** of the HSM.

We will approach these questions one by one in the following subsections.

3.1 Inertial HSM motion

First, there are several ways that we can approach motion. There is periodic, aperiodic and continuous motion. There is also linear motion as well as rotation. We can also vary the degree of electronic control in this motion. The main constraints we have on the HSM's motion pattern are that it needs to be (almost) continuous so as to not expose any weak spots during instantaneous standstill of the HSM. Additionally, for space efficiency the HSM has to stay within a confined space. This means that linear motion must be periodic, like that of a pendulum. 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. In case it has a fixed axis, rotation will expose a weak spot at the axis of rotation where the surface's tangential velocity is low. Faster rotation can lessen the severity of this at the expense of power consumption and mechanical load but can never eliminate it. This effect can be alleviated in two ways: Either by adding additional tamper protection at the axis, or by having the HSM perform a compound rotation that has no fixed axis.

A beneficial side-effect of rotation is that an attacker trying to follow the motion would have to rotate around the same axis. By choosing a suitable rotation frequency we can thus prevent an attacker from following the devices motion since doing so would subject them to impractically large centrifugal forces. Essentially, this limits the approximate maximum size and mass of an attacker based on an assumption on tolerable centrifugal force.

Large centrifugal acceleration at high speeds 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 a *swivel chair attack* (see Section 4.1).

In this paper we focus on rotating IHSMs for simplicity of construction. For our initial research, we are focusing on systems having a fixed axis of rotation due to their relative simplicity in prototype construction but we note the challenge of hardening the shaft against tampering.

3.2 Tamper detection mesh construction

Once we have decided on a type of motion our IHSM's security barrier shall perform, what remains is the actual implementation of that security barrier. There are two movements that we have observed that are key to our work. On the one hand, there is the widespread industry use of delicate tamper sensing mesh membranes. The widespread usage of such membranes in systems deployed in the field for a variety of use cases from low-security payment processing devices to high-security certificate management at a minimum tells us that a properly implemented mesh *can* provide a significant level of security. On the other hand, research has mostly focused on various ways to fabricate enclosures that embed characteristics of a physically uncloneable function (PUF). By using stochastic properties of the enclosure material to form a PUF, such academic designs effectively leverage signal processing techniques to improve the system's hardness by a large factor.

In our research, we focus on security meshes as our IHSM's tamper sensors. Most of the cost in commercial security mesh implementations lies in the advanced manufacturing

techniques and special materials necessary to achieve a sensitive mesh at fine structure sizes. The foundation of an IHSM security is that by moving the mesh even a primitive, coarse mesh made e.g. from mesh traces on a PCB becomes very hard to attack in practice. This allows us to a simple construction made up from low-cost components. Additionally, use of a mesh allows us to only spin the mesh itself around and keep the payload stationary in the center of the IHSM. Tamper sensing technologies that use the entire volume of the HSM such as RF-based systems do not allow for this degree of freedom in their design. They would require the entire IHSM to spin, including its payload. This would entail costly and complex systems for data and power transfer from the outside to the payload.

3.3 Braking detection

The security mesh is a critical component in the IHSM's primary 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 slowing of the IHSM's rotation. Ideally, a sufficiently sensitive sensor should be able to measure any external force applied to the IHSM's rotor and should already trigger a response on the attempt of manipulation.

While the obvious choice to monitor rotation would be a tachometer such as a magnetic or optical sensor attached to the IHSM's shaft, this would be a poor choice in our application. Both optical and magnetic sensors are susceptible to contact-less interference from outside. Another option would be to use feedback from the motor driver electronics. When using a BLDC motor, the driver electronics precisely know the rotor's position at all times. The issue with this approach is that depending on construction, it might invite attacks at the mechanical interface between 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 like a magnetic tachometer or feedback from a BLDC controller, an accelerometer placed inside the rotor's mesh monitoring circuit is a good component to serve as an IHSM's tamper sensor. 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 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 [si2016, 12, 3, 5].

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 target speed of rotation, the accelerometer's location has to be carefully chosen to maximize dynamic range. A key point here is that for rotation speeds between 500 and 1000 rpm, centrifugal acceleration already becomes very large at a radius of just a few cm. At 1000 rpm = 17 Hz at a 10 cm radius acceleration already is above 1000 m s^{-1} or $100 g$. Off-axis performance of commercial accelerometers is usually in the order of 1% so this large acceleration will feed through into all accelerometer axes, even those that are tangential to the rotation. It also means that we either have to place the accelerometer close to the axis or we are limited to a small selection of high- g accelerometers mostly used in automotive applications.

To evaluate the feasibility of accelerometers as tamper sensors we can use a simple benchmark: Let us assume that an IHSM is spinning at 1000 rpm and that we wish to detect any attempt to brake it below 500 rpm. The difference in centrifugal acceleration will be a factor of $\frac{\omega_2^2}{\omega_1^2} = 4$. This results in a factor-4 difference in absolute acceleration




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

that our accelerometer must be able to detect. If we choose our accelerometer’s location to maximize its dynamic range, any commercial MEMS accelerometer should suffice for this degree of accuracy. For rapid deceleration, commercial accelerometers will be much more sensitive as effects of long-term drift can be ignored. If we wish to also detect very slow deceleration, we have to take into account the accelerometer’s drift characteristics.

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

3.4 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 like shown in Figure 2. Shown are the axis of rotation, an accelerometer on the rotating part used to detect braking, the protected payload and the area covered by the rotating tamper detection mesh.

A key observation is that we only have to move the tamper protection mesh, not the entire contents of the HSM. The HSM’s payload and with it most of the HSM’s mass can be stationary. This reduces the moment of inertia of the moving part.

This basic schema accepts a weak spot at the point where the shaft penetrates the spinning mesh. This trade-off makes for a simple mechanical construction and allows power and data connections to the stationary payload through a hollow shaft.

The spinning mesh must be designed to cover the entire surface of the payload, but in contrast to a traditional HSM 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 [10]. This limits the maximum power dissipation of the payload and thus its processing power. Our setup allows direct air cooling of regular heatsinks. This greatly increases the maximum possible power dissipation of the payload and unlocks much more powerful processing capabilities. In an evolution of our design, the spinning mesh could even be designed to *be* a cooling fan.

4 Attacks

After outlining the basic mechanical design of an inertial HSM 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 that is also used in traditional HSMs. This means that in the end an attacker will have to perform the same steps they would have to perform to attack a traditional HSM. Only to attack an IHSM, assuming that the braking detection system works they will have to perform these steps with a tool that follows the HSMs rotation at high speed. This may require specialized mechanical tools, CNC actuators or even a contactless attack using a laser, plasma jet or water jet.

4.1 Swivel chair attacks

First we will consider the most basic of all attacks: A human attacker holding a soldering iron trying to rotate themselves 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 20g as the upper end of the range tolerable by humans for seconds at a time or longer. We thus set our target acceleration to $100g \approx 1000 \text{ m/s}^2$, 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 $f_{\min} = \frac{1}{2\pi} \sqrt{\frac{a}{r}} = \frac{1}{2\pi} \sqrt{\frac{1000 \text{ m/s}^2}{100 \text{ mm}}} \approx 16 \text{ Hz} \approx 1000 \text{ rpm}$. 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.

4.2 Mechanical weak spots


The tamper defense of an IHSM rests on the security mesh moving too fast to tamper. Depending on the type of motion used, the meshes speed may vary by location and over time. Our example configuration of a rotating mesh can keep moving continuously, so it does not have any time-dependent weak spots. It does however have a weak spot at its axis of rotation, at the point where the shaft penetrates the mesh. The meshes 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.

This issue is related to the issue conventional HSMs also face with their power and data connections. In conventional HSMs, power and data are routed into the enclosure through the PCB or flat flex cables sandwiched in between security mesh foil layers. By using a thin substrate and by creating a meandering path by folding the interconnect substrate/security mesh layers several times, in traditional HSMs this interface rarely is a mechanical weak spot. In inertial HSMs, careful engineering is necessary to achieve the same effect. Figure 3 shows variations of the shaft interface of increasing level of complexity.

4.3 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 to allow for a hole to be cut. The other option is to tamper with the monitoring circuit to prevent a damaged mesh from triggering an alarm [15].

Attacks in both locations are electronic attacks, i.e. they require electrical contact to parts of the circuit. Traditionally, this contact is made by soldering or by placing a probe such as a thin needle. We consider this type of attack hard to perform on an object



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

(b) An internal counter-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 3: Mechanical countermeasures to attacks through or close to a rotating IHSM's shaft.

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.

4.4 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.

MEMS accelerometers usually use a cantilever design, where a proof mass moves a cantilever whose precise position is measured electronically. A topic of recent academic interest have been acoustic attacks tampering with these mechanics [21], 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.

4.5 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 this link must be bidirectional so link latency can be measured continuously.

4.6 Fast and violent attacks

A variation of the above attacks on the alarm circuitry is to simply destroy the part of the HSM that erases data in response to tampering before it can perform its job using a tool such as a large hammer or a gun. To mitigate this type of attack, the HSM's tamper response circuitry must be mechanically robust enough to withstand an attack for long enough to carry out its function or else to reliably destroy the payload during an attack.

5 Prototype implementation

As we elaborated above, the mechanical component of an IHSM significantly increases the complexity of any successful 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 prototype IHSM. The main engineering challenges we set out to solve in this prototype were:

1. Fundamental mechanical design suitable for rapid prototyping that can withstand at least 500 rpm.
2. Automatic generation of security mesh PCB layouts for quick adaptation 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 prototype to have space for up to two full-size Raspberry Pi boards for an approximation of 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 ?? shows the rotor's mechanical PCB designs in FreeCAD. The design uses a 6 mm brass tube as its shaft, which is already 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 PCBs (cf. Figure 5b). 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 using a plugin for the KiCAD EDA suite¹. Figure 5a 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 then traced according to a depth-first search through this tree. We consider the quality of the plugin's output sufficient for practical applications. Along with FreeCAD's KiCAD StepUp plugin, this results in an efficient toolchain from mechanical CAD design to production-ready PCB files.

¹<https://blog.jaseg.de/posts/kicad-mesh-plugin/>




(a) The 3D CAD design of the prototype. (b) Assembled mechanical prototype rotor (left) and stator (right) PCB components.

Figure 4: Our prototype IHSM's PCB security mesh design



(a) Overview of the automatic security mesh generation process. 1 - the blob is the example target area. 2 - A grid is overlaid. 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 5: Our automatic security mesh generation process

5.3 Power transmission through the rotating joint

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 have 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 $T_{tx} = 10$ ms. At 100 kBd a transmission of a one-byte message in standard UART framing would take 100 μ s and yield an 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 μ A. Reserving another 100 μ A for the monitoring circuit itself we arrive at an energy consumption of 1.7 A h a⁻¹ (Ampère hour per year).

The annual energy consumption we calculated above is about equivalent to the capacity of a single CR123A lithium primary cell. Using several such cells or optimizing power consumption would thus easily yield several years of battery life. In our prototype we decided against using a battery to reduce rotor mass and 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 μ F ceramic buffer capacitor through a Schottky diode. This solution provides around 3.0 V at several tens of mA to the payload when illumination using either a 60 W incandescent light bulb or a flicker-free LED studio light of similar brightness².

5.4 Data transmission through the rotating joint

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 a MCP6494 general purpose opamp configured as an 100 k Ω transimpedance amplifier. As shown in Figure 6b, 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 6a shows the physical arrangement of both links. The links face opposite one another and are shielded by the motor's body in the center of the PCB.

6 Using MEMS accelerometers for braking detection

Using the prototype from the previous section, 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 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

²LED lights intended for room lighting exhibit significant flicker that can cause the monitoring circuit to reset. Incandescent lighting requires some care in shielding the IR data link from interference.

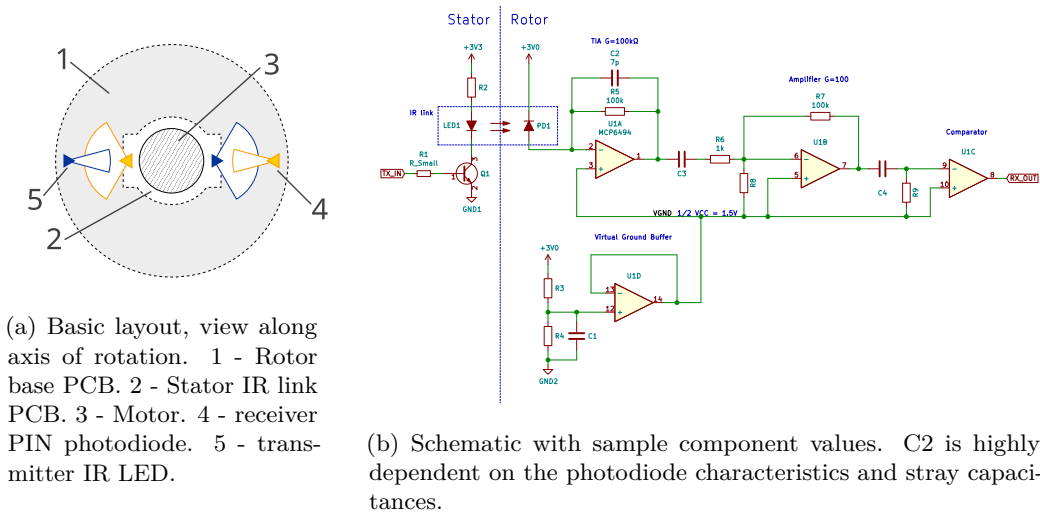


Figure 6: IR data link implementation

tester. We measure the devices rotation speed using a magnet fixed to the rotor and a reed switch held closeby by an articulating arm. The reed switch output is digitized using an USB logic analyzer at a sampling rate of 100 MHz. We calculate rotation frequency as a 1 s running average over debounced interval lengths of this captured signal³.

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 sampling 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 rotors 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’s primary intrinsic errors are offset error and scale error. 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

³A regular frequency counter or commercial tachometer would have been easier, but were not available in our limited COVID-19 home office lab.




Figure 7: Centrifugal acceleration versus angular frequency in theory and in our experiments. Experimental measurements are shown after correction for device-specific offset and scale error. As is evident, our measurements agree very well with our theoretical results. Above 300 rpm, the relative acceleration error was consistently below 0.5%. Below 300 rpm, residual offset error remaining after our first-order corrections has a strong impact (0.05 g absolute or 8% relative at 95 rpm).

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 speeds it causes a large relative error of up to 10% (at 95 rpm).

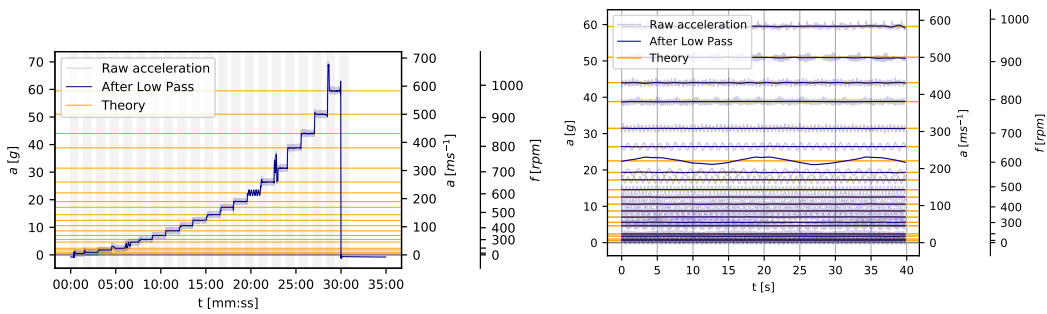
After offset and scale correction, we applied a low-pass filter to our data. The graphs show both raw and filtered data. Raw data contains significant harmonic content. This content is due to vibrations in our prototype. FFT analysis shows that this harmonic content is a clean intermodulation product of the accelerometers sampling rate and the speed of rotation with no other visible artifacts.

Figure 7 shows a plot of our measurement results against frequency. Data points are shown in dark blue, and theoretical behavior is shown in orange.

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 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. Our simple IR data link already is sufficiently reliable for telemetry. Our experiments with the AIS1120 off-the-shelf automotive accelerometer showed that this part is well-suited for braking detection in the range of rotation speed relevant to the IHSM scenario.

Inertial HSMs offer a high level of security beyond what traditional techniques can offer. They allow the construction of devices secure against a wide range of practical attacks at



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

(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 device vibration are clearly visible.

Figure 8: Traces of acceleration measurements during one experiment run.

prototype 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 secure hardware.

References

- [1] 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.
- [2] Ross Anderson. *Security Engineering*. Sept. 16, 2020. ISBN: 978-1-119-64281-7.
- [3] Bertrand Campagnie. *Choose the Right Accelerometer for Predictive Maintenance*. Tech. rep. Analog Devices, 2019.
- [4] 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.
- [5] Maged Elsaid Elnady. “On-Shaft Vibration Measurement Using a MEMS Accelerometer for Faults Diagnosis in Rotating Machines”. PhD thesis. University of Manchester, 2013.
- [6] Jessie Frazelle. “Securing the Boot Process: The hardware root of trust”. In: *ACM Queue* (Dec. 1, 2019). DOI: 10.1145/3380774.3382016.
- [7] Lester Haines. *US outfit patents ‘invisible’ UAV: Stealth through persistence of vision*. Ed. by The Register. Sept. 25, 2006. URL: https://www.theregister.com/2006/09/25/phantom_sentinel/ (visited on 09/17/2020).
- [8] Vincent Immler et al. “Secure Physical Enclosures from Covers with Tamper-Resistance”. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2019). ISSN: 2569-2925. DOI: 10.13154/tches.v2019.i1.51-96.
- [9] International Atomic Energy Agency. *Safeguards, techniques and equipment*. Vol. 1. International Nuclear Verification Series. 2011. ISBN: 978-92-0-118910-3.

- [10] Phil Isaacs et al. *Tamper proof, tamper evident encryption technology*. Tech. rep. Surface Mount Technology Association, 2013.
- [11] Scott Johnson et al. “Titan: enabling a transparent silicon root of trust for Cloud”. In: *Hot Chips: A Symposium on High Performance Chips*. 2018.
- [12] Ivar Koene, Raine Viitala, and Petri Kuosmanen. “Internet of Things Based Monitoring of Large Rotor Vibration With a Microelectromechanical Systems Accelerometer”. In: *IEEE Access* (2019). DOI: <https://doi.org/10.1109/ACCESS.2019.2927793>.
- [13] 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.
- [14] Lily Hay Newman. *Apple’s T2 Security Chip Has an Unfixable Flaw*. Wired Magazine. Oct. 6, 2020. URL: <https://www.wired.com/story/apple-t2-chip-unfixable-flaw-jailbreak-mac/>.
- [15] Karsten Nohl, Fabian Bräunlein, and dexter. *Shopshifting: The potential for payment system abuse*. 32C3 Chaos Communication Congress. Dec. 27, 2015. URL: <https://media.ccc.de/v/32c3-7368-shopshifting#t=2452>.
- [16] 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. ISSN: 2509-3428. DOI: 10.1007/s41635-018-0045-2.
- [17] Mujib Rahman. “Optical fiber cable with tampering detecting means”. US Patent US4859024A. Mar. 10, 1988.
- [18] Sean Smith and Steve Weingart. *Building a High-Performance, Programmable Secure Coprocessor*. Tech. rep. IBM T.J. Watson Research Center, Feb. 19, 1998.
- [19] Daniel Terdiman. *Aboard America’s Doomsday command and control plane*. cnet.com. July 23, 2013. URL: <https://www.cnet.com/news/aboard-americas-doomsday-command-and-control-plane>.
- [20] 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* (Mar. 13, 2020).
- [21] 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.
- [22] Serge Vrijaldenhoven. “Acoustical Physical Uncloneable Functions”. MA thesis. Technische Universiteit Eindhoven, Oct. 1, 2004.

This is version v1.5-eprint-19-g80d1f95-dirty of this paper, generated on April 6, 2021. The git repository can be found at:

<https://git.jaseg.de/rotohsn.git>