

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

Jan Götte

ABSTRACT

In this paper, we introduce a novel countermeasure against physical attacks: Inertial hardware security modules. 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.

I. 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 [10, 4, 8].

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 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 [3]. 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 Appendix B). 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.

In Section II, we will give an overview of the state of the art in the physical security of HSMs. On this basis, in Section III we will elaborate the principles of our inertial HSM approach. We will analyze its weaknesses in Section IV. Based on these results we have built a prototype system that we will illustrate in Section V. We conclude this paper with a general evaluation of our design in Section VI.

II. RELATED WORK

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 [7, 6, 2]. There has been some research on monitoring the HSM's inside using e.g. electromagnetic radiation [16, 9] or ultrasound [18] but none of this research has found widespread adoption.

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 [14]. This HSM is an example of an industry-standard construction. Though 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. Apart from 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 [14], 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 [12, 3, 2, 7].

In [6], 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 [16], 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 [16] 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 [6], at the cost of worse and less predictable security guarantees. Where [16] use electromagnetic radiation, Vrijaldenhoven in [18] uses ultrasound waves travelling on a surface acoustic wave (SAW) device to a similar end.

While [16] approach the sensing frontend cost as their only optimization target, the prior work of Kreft and Adi [9] 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.

Our concept is novel in that mechanical motion has not been proposed before as part of a hardware security module. Most academic research concentrates on the issue of creating new,




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

more sensitive security barriers for HSMs [6] while commercial vendors concentrate on means to cheaply manufacture and certify these security barriers [3]. 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 very high-performance one. The closest to a mechanical HSM that we were able to find during our research is an 1988 patent [13] that describes a mechanism to detect tampering along a communication cable by enclosing the cable inside a conduit filled with pressurized gas.

III. INERTIAL HSM CONSTRUCTION AND OPERATION

Mechanical motion has been proposed as a means of making things harder to see with the human eye [5] and is routinely used in military applications to make things harder to hit [15] but we seem to be the first to use it in tamper detection. If we consider different ways of moving an HSM to make it harder to tamper with, we find that making it spin has several advantages.

First, the HSM has to move fairly fast. If any point of the HSM's tamper sensing mechanism moves slow enough for a human to follow, it becomes a weak spot. E.g. in a linear pendulum motion, the pendulum becomes stationary at its apex. Second, a spinning HSM is compact compared to alternatives like an HSM on wheels. Finally, rotation leads to predictable accelerometer measurements. A beneficial side-effect of spinning the HSM is that if the axis of rotation is within the HSM itself, an attacker trying to follow the motion would have to rotate around the same axis. Their tangential linear velocity would rise linearly with the radius from the axis of rotation, which allows us to limit the approximate maximum size and mass of an attacker using an assumption on tolerable centrifugal force (see Appendix B). In this consideration the axis of rotation is a weak spot, but that can be mitigated using multiple nested layers of protection.

In a rotating reference frame, centrifugal force is proportional to the square of angular velocity and proportional to distance from the axis of rotation. We can exploit this fact to create a sensor that detects any disturbance of the rotation by placing a linear accelerometer at some distance from the axis of rotation. During constant rotation, both acceleration

tangential to the rotation and along the axis of rotation will be zero. Centrifugal acceleration will be constant.

Large centrifugal acceleration at high speeds poses the engineering challenge of preventing the whole thing from flying apart, but it also creates an obstacle to any attacker trying to manipulate the sensor. We do not need to move the entire contents of the HSM. It suffices if we move the tamper detection barrier around a stationary payload. This reduces the moment of inertia of the moving part and it means we can use cables for payload power and data.

From our back-of-the-envelope calculation in Appendix B we conclude that even at moderate speeds above 500 rpm, an attack would have to be carried out using a robot.

In Appendix C we consider sensor configurations and we conclude that one three-axis accelerometer each in the rotor and in the stator are a good baseline configuration. In general, the system will be more sensitive to attacks if we over-determine the system of equations describing its motion by using more sensors than necessary.

A. Mechanical layout

The simplest way to mount a stationary payload in a spinning security mesh is to use a hollow shaft. The payload can be mounted on a fixed rod threaded through this hollow shaft along with wires for power and data. The shaft is a weak spot of the system, but this weak spot can be alleviated through either careful construction or a second layer of rotating meshes with a different axis of rotation. Configurations that do not use a hollow-shaft motor are possible, but may require additional bearings to keep the stator from vibrating.

The spinning mesh must be designed to cover the entire surface of the payload during one revolution. Still, it can be designed with longitudinal gaps to allow outside air to flow through to the payload. In boundary-sensing HSMs, cooling of the processor inside 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 [7]. This limits the maximum power dissipation of the payload and thus its processing power. Our setup allows direct air cooling, which increases the maximum possible power dissipation of the payload and unlocks much more powerful processing capabilities. Instead of gaps one could even integrate an actual fan into the rotor.

B. Spinning mesh power and data transmission

The basic concept of a security mesh spinning at more than 500 rpm around a payload leaves us with a few implementation challenges. Since the spinning mesh must be monitored for breaks or short circuits continuously, we need both a power supply for the spinning monitoring circuit and a data link back to the stator.

A good starting point for power transfer is a simple setup of a stationary bright lamp shining at a rotating solar panel. In contrast to e.g. slip rings, this setup is mechanically durable at high speeds and it also provides reasonable output power (see Appendix A for some calculations on power consumption).

A battery may not provide a useful lifetime without power-optimization. Likewise, an energy harvesting setup may not provide enough current to supply peak demand.

Since the monitoring circuit uses little current, power transfer efficiency is not important. On the other hand, cost may be a concern in a production device. Here it may prove worthwhile to replace the solar cell setup with an extra windings on the rotor of the BLDC motor driving the spinning mesh. This rotor is likely to be a custom part, so adding these windings is unlikely to increase cost significantly. Inductive power transfer may also be an option given that one can integrate it into the mechanical design.

Besides power, the data link between spinning mesh and payload is critical to the HSM's design. This link is used to transmit the occasional status report along with a low-latency alarm trigger ("heartbeat") signal from mesh to payload. As we will elaborate in Section V a simple infrared optical link turned out to be a good solution for this purpose.

IV. ATTACKS

After outlining the basic mechanical design of an inertial HSM above, in this section we will detail possible ways to attack it. Fundamentally, attacks on an inertial HSM are the same as those on a traditional HSM, since the tamper detection mesh is the same. Only in the inertial HSM any attack on the mesh has to be carried out while the mesh is rotating, which for most types of attack will require a CNC attack robot moving in sync with it. In comparison to traditional designs, the data link between mesh and payload is an additional weak spot in the rotating desing. If it is optical, non-contact attacks are possible.

A. Attacks on the mesh

There are two locations where one can attack a tamper-detection mesh. On one hand, the mesh itself can be tampered with. This includes bridging its traces to allow for a hole to be cut. The other option is to tamper with the monitoring circuit itself, to prevent a damaged mesh from triggering an alarm and causing the HSM to erase its contents [11]. 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 contact infeasible to be performed on an object spinning at high speed without a complex setup that rotates along with the object or that involves ion beams, electron beams or liquids. Thus, we consider them to be practically infeasible outside of a well-funded, special-purpose laboratory.

B. Attacks on the alarm circuitry

An electronic attack could also target the alarm circuitry inside the stationary payload, or the communication link between rotor and payload. The link can easily be proofed by using a cryptographically secured protocol 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 and has to tolerate environmental attacks such as ones using temperature, ionizing radiation, lasers, supply voltage variations, ultrasound or other vibration and gases or liquids. The easiest way to proof an alarm system against these is to employ adequate filtering of the incoming power supply and use sensors for the others, triggering an alarm in case extraordinary environmental variations are detected.

If the alarm link between rotor and stator uses a spoofable interface such as an optical link, this link must be cryptographically verified. It also must be bidirectional to allow the alarm signal receiver to verify link latency. In a purely unidirectional spoofable link, an attacker could record the authenticated “no alarm” signal from the transmitter while simultaneously replaying it just slightly slower (say at 99% speed) to the receiver. The receiver would not be able to distinguish between this attack and ordinary deviations in the transmitter’s local clock frequency. However, the attacker can at any point simply stop the rotor and replay the leftover recorded “no alarm” signal. Given the frequency stability of commercial crystals, this would allow for an attack duration of several seconds per hour of recording time.

C. Fast and violent attacks

A variation of the above attacks on the alarm circuitry would be an attack that attempts to simply destroy this circuitry before the alarm can be acted upon using a tool like a large hammer or a gun. Mitigations for this type of attack include potting the payload inside a mechanically robust enclosure. The alarm signalling chain’s integrity can be checked continuously using a cryptographic heartbeat protocol. A simple active-high or active-low alarm signal cannot be considered fail-safe in this scenario.

D. Attacks on the rotation sensor

An attacker may try to stop the rotor before tampering with the mesh. To succeed, they would need to fool the rotor’s MEMS accelerometer. An electronic attack on the sensor or the monitoring microcontroller would be no easier than directly bridging the mesh traces. Physical attacks on the accelerometer are possible [17], but in the authors’ estimate are too hard to control to be practically useful.

A last type of attack might be to try to physically tamper with the accelerometer’s sensing mechanism. MEMS accelerometers usually use a cantilever design, where a proof mass moves a cantilever whose precise position can be measured electronically. A possible way to attack such a device might be to first decapsulate it using laser ablation synchronized with the device’s rotation. Then, a fast-setting glue such as a cyanoacrylate could be deposited on the moving MEMS parts, locking them in place. This attack would require direct access to the accelerometer from the outside and can be prevented by mounting the accelerometer in a shielded place inside the security envelope. This attack can only work if the rate of rotation and thus the accelerometer’s readings are constant. If the rate of rotation is set to change on a schedule, this type of attack can be detected easily. In Appendix C we outline the constraints on sensor placement.




Figure 2: The 3D CAD design of the prototype.

V. PROTOTYPE IMPLEMENTATION

To validate our theoretical design, we implemented a prototype rotary HSM. The main engineering challenges we solved in our prototype are:

- 1) Fundamental mechanical design suitable for rapid prototyping that can withstand a rotation of 500 rpm.
- 2) Automatic generation of security mesh PCB layouts for quick adaption to new form factors.
- 3) Non-contact power transmission to rotor.
- 4) Non-contact bidirectional data communication between stator and rotor.

A. Mechanical design

We sized our prototype to have space for up to two full-size Raspberry Pi boards. Each one of these boards is already more powerful than an ordinary HSM, but they are small enough to simplify our prototype’s design. For low-cost prototyping we designed our prototype to use printed circuit boards as its main structural material. The interlocking parts were designed in FreeCAD as shown in Figure 2. The mechanical designs were exported to KiCAD for electrical design before being sent to a commercial PCB manufacturer. Rotor and stator are built from interlocking, soldered PCBs. The components are mounted to a 6 mm brass tube using FDM 3D printed flanges. The rotor is driven by a small hobby quadcopter motor.

Security is provided by a PCB security mesh enveloping the entire system and extending to within a few millimeters of the shaft. For security it is not necessary to cover the entire circumference of the module with mesh, so we opted to use only three narrow longitudinal struts to save weight.

To mount the entire HSM, we chose to use “2020” modular aluminium profile.

B. PCB security mesh generation

To allow a quick iteration of our design while producing results with a realistic level of security, we wrote a plugin for the




Figure 3: 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.




Figure 4: A section of the security mesh PCB we produced with our toolchain for the prototype HSM.

KiCAD EDA suite that automatically generates parametrized security meshes. When KiCAD is used in conjunction with FreeCAD through FreeCAD's KiCAD StepUp plugin, this ends up in an efficient toolchain from mechanical CAD design to security mesh PCB gerber files. The mesh generation plugin can be found at its website¹.

Our mesh generation plugin overlays a grid on the target area and then produces a randomized tree covering this grid. The individual mesh traces are then traced along a depth-first search through this tree. A visualization of the steps is shown in Figure 3. A sample of the production results from our prototype is shown in Figure 4.

C. Data transmission through rotating joint

As a baseline solution for data transmission, we settled on a 115 kBd UART signal sent through a simple bidirectional infrared link. In the transmitter, the UART TX line on-off modulates a 920 nm IR LED through a common-emitter driver transistor. In the receiver, an IR PIN photodiode reverse-biased to $\frac{1}{2}V_{CC}$ is connected to a reasonably wideband transimpedance amplifier (TIA) with a $100\text{ k}\Omega$ transimpedance. As shown in Figure 6, the output of this TIA is fed through another $G = 100$ amplifier whose output is then squared up by a comparator. We used an MCP6494 quad CMOS op-amp. At a specified 2 mA current consumption it is within our rotor's power budget, and its Gain Bandwidth Product of 7.5 MHz

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




Figure 5: Schema of our bidirectional IR communication link between rotor and stator, view along axis of rotation. 1 - Rotor base PCB. 2 - Stator IR link PCB. 3 - Motor. 4 - receiver PIN photodiode. 5 - transmitter IR LED.

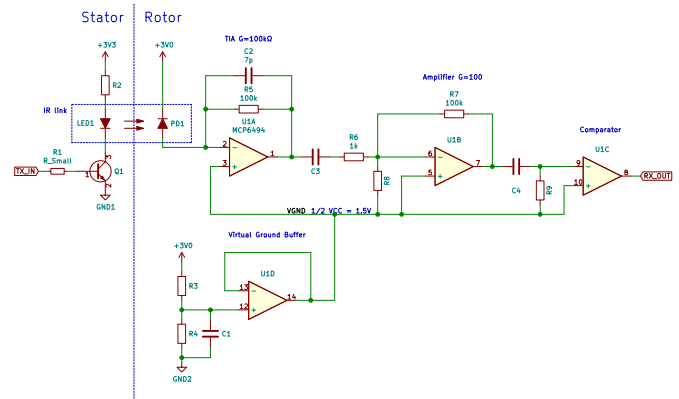


Figure 6: Schematic of the IR communication link. Component values are only examples. In particular $C2$ depends highly on the photodiode used and stray capacitances due to the component layout.

yields a useful transimpedance in the photodiode-facing TIA stage.

To reduce the requirements on power transmission to the rotor, we have tried to reduce power consumption of the rotor-side receiver/transmitter pair trading off stator-side power consumption. One part of this is that we use a wide-angle photodiode and IR LED on the stator, but use narrow-angle components on the rotor. The two rx/tx pairs are arranged next to the motor on opposite sides. By placing the narrow-angle rotor rx/tx components on the outside as shown in Figure 5, the motor shields both IR links from crosstalk. The rotor transmitter LED is driven at 1 mA while the stator transmitter LED is driven at 20 mA.

D. Power transmission through rotating joint

Since this prototype serves only demonstration purposes, we chose to use the simplest possible method of power transmission: solar cells. We mounted six series-connected solar cells in three commercially available modules on the circular PCB at the end of our cylindrical rotor. The solar cells directly feed the rotor's logic supply with buffering by a large $33\text{ }\mu\text{F}$ ceramic capacitor. With six cells in series, they




Figure 7: The prototype when we first achieved reliable power transfer and bidirectional communication between stator and rotor. In the picture, the prototype was communicating reliably up to the maximum ≈ 1500 rpm that we could get out of its hobby quadcopter parts.

provide around 3.0 V at several tens of mA given sufficient illumination.

For simplicity and weight reduction, at this point we chose to forego large buffer capacitors on the rotor. This means variations in solar cell illumination directly couple into the microcontroller’s supply rail. Initially, we experimented with regular residential LED light bulbs, but those turned out to have too much flicker and lead to our microcontroller frequently rebooting. Trials using an incandescent light produced a stable supply, but the large amount of infrared light emitted by the incandescent light bulb severely disturbed our near-infrared communication link. As a consequence of this, we settled on a small LED light intended for use as a studio light that provided us with almost flicker-free light at lower frequencies, leading to a sufficiently stable microcontroller VCC rail without any disturbance to the IR link.

E. Evaluation

During experiments, our prototype performed as intended. After some experimentation, we got both power and data transmission through the rotating joint working reliably. Figure 7 shows our prototype performing reliably at maximum speed for the first time. Our improvised IR link is open in both directions for about 60° of the rotation, which allows us to reliably transfer several tens of bytes in each direction during each receiver’s fly-by even at high speed of rotation. As a result of our prototype experiments, we consider a larger-scale implementation of the inertial HSM concept practical.

VI. CONCLUSION

In this paper, we introduced inertial hardware security modules (iHSMs), a novel concept for the construction of highly secure hardware security modules from inexpensive, commonly available parts. We elaborated the engineering considerations underlying a practical implementation of this

concept. We implemented a prototype demonstrating practical solutions to the significant engineering challenges of this concept. We analyzed the concept for its security properties and highlighted its ability to significantly strengthen otherwise weak tamper detection barriers.

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 prototype quantities and without specialized tools. 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] 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.
- [4] Jessie Frazelle. “Securing the Boot Process: The hardware root of trust”. In: *ACM Queue* (Dec. 1, 2019). DOI: 10.1145/3380774.3382016.
- [5] 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).
- [6] 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.
- [7] Phil Isaacs et al. *Tamper proof, tamper evident encryption technology*. Tech. rep. Surface Mount Technology Association, 2013.
- [8] Scott Johnson et al. “Titan: enabling a transparent silicon root of trust for Cloud”. In: *Hot Chips: A Symposium on High Performance Chips*. 2018.
- [9] 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.
- [10] 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/>.
- [11] 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>.

- [12] 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.
- [13] Mujib Rahman. “Optical fiber cable with tampering detecting means”. US Patent US4859024A. Mar. 10, 1988.
- [14] Sean Smith and Steve Weingart. *Building a High-Performance, Programmable Secure Coprocessor*. Tech. rep. IBM T.J. Watson Research Center, Feb. 19, 1998.
- [15] Daniel Terdiman. *Aboard America’s Doomsday command and control plane*. cnet.com. July 23, 2013. URL: <https://www.cnet.com/news/aboard-america-doomsday-command-and-control-plane>.
- [16] 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).
- [17] 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.
- [18] Serge Vrijaldenhoven. “Acoustical Physical Uncloenable Functions”. MA thesis. Technische Universiteit Eindhoven, Oct. 1, 2004.

APPENDIX

A. Spinning mesh energy calculations

Assume 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 $100\ \mu\text{s}$ 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 $100\ \mu\text{A}$. Reserving another $100\ \mu\text{A}$ for the monitoring circuit itself we arrive at an energy consumption of $1.7\ \text{A h a}^{-1}$.

1) *Battery power*: 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.

2) *LED and solar cell*: Let us assume an LED with a light output of 1 W illuminating a small solar cell. Let us pessimistically assume a 5% conversion efficiency in the solar cell. Let us assume that when the rotor is at its optimal rotational angle, 20% of the LED’s light output couple into the solar cell. Let us assume that we lose another 90% of light output on average during one rotation when the rotor is in motion. This results in an energy output from the solar cell of 1 mW. Assuming a 3.3 V supply this yields $300\ \mu\text{A}$ for our monitoring circuit. This is enough even with some conversion losses in the step-up converter boosting the solar cell’s 0.6 V working voltage to the monitoring circuit’s supply voltage.

B. Minimum angular velocity: Rotating human attacker

An attacker might try to rotate along with the HSM to attack the security mesh without triggering the accelerometer. Let us pessimistically assume that the attacker has the axis of rotation running through their center of mass. The attacker’s body is probably at least 200 mm wide along its shortest axis, resulting in a minimum radius from axis of rotation to surface of about 100 mm. We choose $250\ \text{m/s}^2$ as an arbitrary acceleration well past the range tolerable by humans according to Wikipedia. Centrifugal acceleration is $a = \omega^2 r$. In our example this results in a minimum angular velocity of $\omega_{\min} = \sqrt{\frac{a}{r}} = \sqrt{\frac{250\ \text{m/s}^2}{100\ \text{mm}}} \approx 8 \cdot 2\pi \frac{1}{\text{s}} \approx 500\ \text{rpm}$.

C. Fooling the accelerometer

Let us consider a general inertial HSM with one or more sensors that is attacked by an attacker. In this scenario, it is reasonable to assume that the rotating parts of the HSM are rigidly coupled to one another and will stay that way: For the attacker to decouple parts of the HSM (e.g. to remove one of its accelerometers from the PCB), the attacker would already have to circumvent the rotor’s security mesh.

Assuming the HSM is stationary, a sensor on the rotating part will experience two significant accelerations:

- 1) Gravity $g = 9.8 \frac{\text{m}}{\text{s}^2}$
- 2) Centrifugal force $a_C = \omega^2 r$, in the order of $1000\ \text{m/s}^2$ or $100g$ at $r = 100\ \text{mm}$ and $1000\ \text{rpm}$

Due to the vast differences in both radius and angular velocity, we can neglect any influence of the earth’s rotation on our system.

In normal operation, the HSM is stationary ($\mathbf{v} = 0$) and the HSM’s motor is tuned to exactly counter-balance friction so the rotor’s angular velocity remains constant. As a rigid body, the rotor’s motion is fully defined by its rotation and translation. In total, this makes for six degrees of freedom. The three degrees of freedom of linear translation we can measure directly with an accelerometer in the stationary part on the inside of the HSM. This accelerometer could detect any rapid acceleration of the HSM’s rotor. To measure rotation, we could mount a gyroscope on the rotor to detect deceleration. The issue with this is that like other MEMS acceleration sensors, commercial MEMS gyroscopes are vulnerable to drift and an attacker could slowly decelerate the rotor without being detected.

A linear accelerometer mounted on the rotor however is able to catch even this attack. Subtracting gravity, it could determine both magnitude and direction of the centrifugal force, which is proportional to the square of angular velocity and not its derivative.

In summary, a single three-axis accelerometer on the rotor combined with a three-axis accelerometer in the stator would be a good baseline configuration.

D. Patents and licensing

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. This paper itself is licensed CC-BY-SA (see below). As for the inertial HSM concept, 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.



This work is licensed under a Creative-Commons “Attribution-ShareAlike 4.0 International” license. The full text of the license can be found at:

<https://creativecommons.org/licenses/by-sa/4.0/>

For alternative licensing options, source files, questions or comments please contact the authors.

This is version v1.3-2-gdabb030a-dirty generated on December 18, 2020. The git repository can be found at:

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