

Can't touch this: Inertial HSMs Foil Advanced Physical Attacks

Jan Götte

ABSTRACT

In this paper, we introduce a novel, highly effective 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 even the best commercial offerings. 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 equates 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 in form of smartcard-like 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 be reduced to that of its physically secured TPM[8, 4, 9]. Being physically small, physical security is less of a challenge on the scale of a TPM.

A. Technical approaches to physical security

Shrinking things to the nanoscopic level to secure them against tampering is an engineering solution to problems that cannot be solved (yet) with cryptographic security. The security of chips like smartcards and TPMs often rests on the assumption that their fine structures are hard to reverse

engineer and modify. As of now, this property holds and in the authors' opinion it will likely be a reasonable assumption 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 these chips[1, 2].

B. Hardware Security Modules

Right now, Hardware security modules (HSMs) are the class of commercial devices offering the highest “physical security-to-volume-product”. Where smartcards physically secure a single chip, HSMs secure a small circuit board. In contrast to a smartcard, in a tradeoff between security and convenience the HSM actively deletes its secrets when it detects a manipulation. Commercial HSMs commonly employ what we call *boundary monitoring*. They have a physical security barrier that they continuously monitor for holes. Usually, this barrier is a thin foil that is patterned with at least two meandering electrical traces that is folded in layers to cover the entire area of the foil. The HSM monitors these traces for shorts or breaks. This simple construction transforms the security problem into a manufacturing challenge[7, 6, 2].

In our classification the other type of HSMs are *volumetric* HSMs. They monitor their entire internal volume for changes using e.g. electromagnetic radiation[17, 11] or ultrasound[19]. Their security is limited by the analog sensitivity of their transceivers. Their practicality is limited by their complex transceiver and signal processing circuitry. They promise to secure larger volumes than boundary monitoring at higher parts cost. A problem with volumetric designs is their security analysis, which is hard to do without significant guesswork. In e.g. a device that use electromagnetic radiation to monitor its volume, one might have to numerically solve the electromagnetic field equations inside the HSM to validate its impenetrability.

C. Inertial HSMs: A new approach to physical security

We are certain that there is still much work to be done and many insights to be gained in both HSM and in smartcard technology¹. Still, we wish to introduce a novel approach to sidestep the issues of conventional HSMs and provide radically better security against physical attacks. Our core observation

¹As a baseline, consider a box with mirrored walls that contains a smaller box suspended on thin wires that has cameras looking outward in all directions at the mirrored walls. Given that the defender can control lighting conditions inside this kaleidoscopic box in this application modern cameras perform better than the human eye. Thus, a successful physical attack on this system would likely an “invisibility cloak”—and the system would remain secure as long as no such thing exists. To be viable, an HSM technology must be either cheaper, smaller or more sensitive than this strawman setup[10].

is that any cheap but coarse HSM technology can be made much more difficult to attack by moving it very quickly. As a trivial example, consider an HSM as it is used in ecommerce applications for credit card payments. 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, triggering 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 life². 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.

D. Contributions

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.

II. RELATED WORK

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 [15]. 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 [15], 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[13, 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 [17], 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 [17] 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 [17] use electromagnetic radiation, Vrijaldenhoven in [19] uses ultrasound waves travelling on a surface acoustic wave (SAW) device to a similar end.

While [17] approach the sensing frontend cost as their only optimization target, the prior work of Kreft and Adi [11] 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, 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 [14] 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

A. Using motion for tamper detection

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[16] but we seem to be the first to use it in tamper detection. Let us think about the constraints of our approach.

- 1) We need the tamper sensor's motion to be fairly fast. If any point of the sensor moves slow enough for a human to follow, it becomes a weak spot.
- 2) We need to keep the entire apparatus compact.

²See Appendix B




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

- 3) We need the sensor's motion to be very predictable so that we can detect an attacker trying to stop it.

From this, we can make a few observations.

- 1) Non-periodic linear motion (like a train on wheels) is likely to be a poor choice since it requires a large amount of space, and it is comparatively easy to follow something moving linearly.
- 2) Oscillatory motion such as linear vibration or a pendulum motion might be a good candidate would there not be the moment at its apex when the vibration reverses direction the object is stationary. This is a weak spot.
- 3) Rotation is a very good choice. It does not require much space to execute. Additionally, 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. Since their tangential linear velocity would rise linearly with the radius from the axis of rotation, an assumption on tolerable centrifugal force allows one to limit the approximate maximum size and mass of an attacker (see Appendix B). The axis of rotation is a weak spot, but we can simply nest multiple layers of protection at an angle to each other.
- 4) We do not have 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.

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. At high speeds, this acceleration may become very large. This poses the engineering challenge of preventing the whole thing from flying apart, but also creates an obstacle to any attacker trying to manipulate the sensor.

In Appendix B we present some back-of-the-envelope cal-

culations on minimum angular velocity. 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. Other configurations such as one using two two-axis accelerometers in the rotor are also possible. 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.

B. Payload mounting mechanisms

The simplest way to mount a stationary payload in a spinning security mesh is to drive the rotor using a hollow shaft. This allows the payload to be mounted on a fixed rod threaded through this hollow shaft along with wires for power and data. The stationary rod and cables on the axis of rotation inside the hollow shaft are 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.

C. Spinning mesh power supply

There are several options to transfer power to the rotor from its stationary frame.

- 1) Slip ring contacts are a poor candidate as they are limited in their maximum speed and lifetime, and as precision mechanical components are expensive.
- 2) Inductive power transfer as used in inductive charging systems can be used without modification if both coils are mounted axially.
- 3) A second brushless motor on the axis of rotation can be used as a generator, with its axis connected to the fixed frame and its stator mounted and connected to the rotor. Likewise, a custom-made drive motor that includes some auxiliary rotor windings for power transfer in addition to the rotor's magnets would be possible.
- 4) A bright lamp along with some small solar cells may be a practical approach for small amounts of energy³.
- 5) For a very low-power security mesh, a battery specified to last for the lifetime of the device may be practical⁴.

In our prototype, we settled on a solar cell-based solution for its simplicity.

D. Payload cooling

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 be solved with complex and costly siphon-style constructions, but in commercial systems heat conduction is used exclusively[7]. This limits the maximum power dissipation of the payload and

³See Appendix A for a back-of-the-envelope calculation

⁴See Appendix A

thus its processing power. In our spinning HSM concept, the spinning mesh can have longitudinal gaps in the mesh without impeding its function. This allows air to pass through the mesh during rotation, and one could even integrate an actual fan into the rotor. This greatly increases the maximum possible power dissipation of the payload and unlocks much more powerful processing capabilities.

E. Spinning mesh data communication

As for power, slip rings are the obvious choice to couple data signals through the rotating joint. Like for power, ones that match our reliability and speed constraints are expensive.

Our design has a stationary payload and only the security mesh and sensors are spinning. The rotor only needs to send occasional status reports and a high-frequency alarm trigger heartbeat signal to the stator. For this, a simple optocoupler close to the axis of rotation is a good solution that we implemented in our prototype.

IV. ATTACKS

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[12]. 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 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[18], 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, it is trivially detectable.

In Appendix C we outline the constraints on sensor placement.

V. PROTOTYPE IMPLEMENTATION

To validate our theoretical design, we have 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.




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

A. Mechanical design

We sized our prototype to have space for one or 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 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 website⁵.

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.

⁵<https://blog.jaseg.de/posts/kicad-mesh-plugin/>




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.

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 k Ω 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 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.




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.

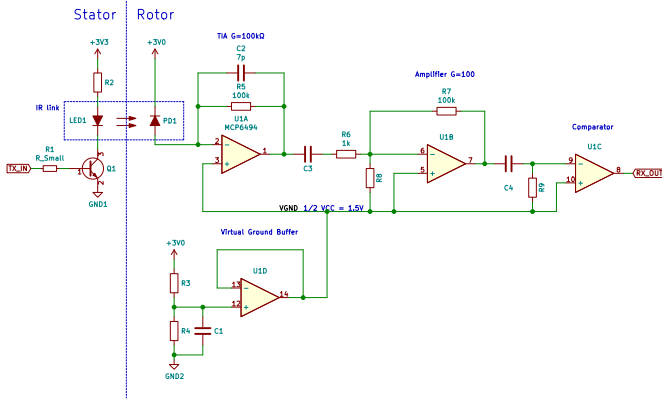


Figure 6: Schematic of the IR communication link. Component values are only examples. In particular C_2 depends highly on the photodiode used and stray capacitances due to the component layout.

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 made up from 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 $33\ \mu\text{F}$ ceramic capacitor. With six cells in series, they provide around $3.0\ \text{V}$ 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




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 $\approx 1500\ \text{rpm}$ that we could get out of its hobby quadcopter parts.

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. FUTURE WORK

A. Design space exploration

There are several aspects of inertial HSM design that we wish to explore in future work.

a) *Other modes of movement:* An oscillating iHSM might enable power and data transfer to the moving part using cables.

b) *Multiple axes of rotation:* The weak spot of our prototype design at the stationary shaft can be alleviated using gyroscope mechanics.

c) *Other sensing modes:* By printing the inside of the rotor with a pattern that is observed by a linear CCD a completely passive rotor may be possible.

d) *Bearing longevity:*

e) *Handling of gyroscopic precession forces during shipping:*

B. Penetration testing

We intend to refine our prototype design to production quality. As part of this, we wish to try out a range of attacks on our prototype.

VII. CONCLUSION

In this paper, we have presented inertial hardware security modules (iHSMs), a novel concept for the construction of highly secure hardware security modules from inexpensive, commonly available parts. We have elaborated the engineering considerations underlying a practical implementation of this concept. We have implemented a prototype demonstrating practical solutions to the significant engineering challenges of this concept. We have analyzed the concept for its security properties and highlighted its ability to significantly strengthen otherwise weak tamper detection barriers. We have laid out some ideas for future research on the concept.

REFERENCES

- [1] Nils Albartus et al. “DANA Universal Dataflow Analysis for Gate-Level Netlist Reverse Engineering”. In: *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2020.4 (2020), pp. 309–336. DOI: 10.13154/tches.v2020.i4.309-336.
- [2] Ross Anderson. *Security Engineering*. Sept. 16, 2020.
- [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”. In: *Pan Pacific Symposium*. 2013.
- [8] *Jailbreaker nehmen T2-Sicherheitschip von Macs ins Visier*.
- [9] Scott Johnson et al. “Titan: enabling a transparent silicon root of trust for Cloud”. In: *Hot Chips: A Symposium on High Performance Chips*. 2018.
- [10] Seung Hyun Kim, Su Chang Lim, et al. “Intelligent intrusion detection system featuring a virtual fence, active intruder detection, classification, tracking, and action recognition”. In: *Annals of Nuclear Energy* 112 (2018), pp. 845–855.
- [11] 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.
- [12] Karsten Nohl, Fabian Bräunlein, and dexter. “Shopshifting: The potential for payment system abuse”. Dec. 27, 2015.
- [13] 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.
- [14] Mujib Rahman. “Optical fiber cable with tampering detecting means”. US4859024A. Mar. 10, 1988.
- [15] Sean Smith and Steve Weingart. *Building a High-Performance, Programmable Secure Coprocessor*. Tech. rep. IBM T.J. Watson Research Center, Feb. 19, 1998.
- [16] Daniel Terdiman. *Aboard America’s Doomsday command and control plane*. July 2013.
- [17] 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).
- [18] 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 (EuroS&P)*. IEEE. 2017, pp. 3–18.
- [19] Serge Vrijaldenhoven. “Acoustical Physical Unclonable 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 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 m/s^2 or $100g$ at $r = 100 \text{ mm}$ and 1000 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 `v1.2-0-g0e6fbee-dirty` generated on December 2, 2020. The git repository can be found at:

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