Ripples in the Pond: Transmitting Information through Grid Frequency Modulation

Jan Sebastian Götte, Liran Katzir, and Björn Scheuermann

[1] HIIG
safetyreset@jaseg.de
[2] Tel Aviv University
Faculty of Engineering
[3] HU Berlin
scheuermann@informatik.hu-berlin.de

Keywords: Security, privacy and resilience in critical infrastructures · Security and privacy in "internet of things" · Cyber-physical systems · Hardware security · Network Security · Energy systems · Signal theory

Abstract. The smart grid is a large, complex and interconnected technological system. With remotely controllable load switches having been rolled out at scale in some countries, a tiny flaw inside the firmware of one of these embedded devices may enable attacks to remotely trigger large-scale disruption with potentially catastrophic results. Attaining perfect security against such cyberphysical attacks is a monumental embedded engineering task—and observations do not indicate that current efforts meet the requirements of this task.

In this paper, we approach the smart grid safety issue by implementing an emergency override that can be used to reset all connected devices to a known-good state and preempt subsequent compromise by cutting communication links. To yield a fully fail-safe design, our system does not rely on the internet or other conventional communication network to work. Instead, our system transmits error-corrected and cryptographically secured commands by modulating grid frequency using a single large consumer such as a large aluminium smelter. This approach differs from traditional Powerline Communication (PLC) systems in that reaches every device within the same synchronous area as the signal is embedded into the fundamental grid frequency instead of a superimposed voltage that is quickly attenuated across long distances.

Using simulations we have determined that control of a 25 MW load would allow for the transmission of a crytographically secured *reset* signal within 15 minutes. We have produced a proof-of-concept prototype receiver that demonstrates the feasibility of decoding such signals even on resource-constrained microcontroller hardware.

1 Introduction

In the power grid, as in many other engineered systems, we can observe an ongoing diffusion of information systems into the domain of industrial control.

Automation of these control systems has already been practiced for the better part of a century. Throughout the 20th century this automation was mostly limited to core components of the grid. Generators in power stations are computer-controlled according to electromechanical and economic models. Switching in substations is automated to allow for fast failure recovery. Human operators are still vital to these systems, but their tasks have shifted from pure operation to engineering, maintenance and surveillance[9, 2].

With the turn of the century came a large-scale trend in power systems to move from a model of centralized generation, built around massive large-scale fossil and nuclear power plants, towards a more heterogenous model of smaller-scale generators working together. In this new model large-scale fossil power plants still serve a major role, but new factors come into play. One such factor is the advance of renewable energies. The large-scale use of wind and solar power in particular seems unavoidable for continued human life on this planet. For the electrical grid these systems constitute a significant challenge. Fossil-fueled power plants can be controlled in a precise and quick way to match energy consumption. This tracking of consumption with production is vital to the stability of the grid. Renewable energies such as wind and solar power do not provide the same degree of controllability, and they introduce a larger degree of uncertainty due to the unpredictability of the forces of nature[9].

Along with this change in dynamic behavior, renewable energies have brought forth the advance of distributed generation. In distributed generation end-customers that previously only consumed energy have started to feed energy into the grid from small solar installations on their property. Distributed generation is a chance for customers to gain autonomy and shift from a purely passive role to being active participants of the electricity market[9].

To match this new landscape unpredictable renewable resources and of decentralized generation, the utility industry has had to adapt itself in major ways. One aspect of this adaptation that is particularly visible to energy consumers is the computerization of end-user energy metering. Despite the widespread use of industrial control systems inside the electrical grid and the far-reaching diffusion of computers into people's everyday lives, the energy meter has long been one of the last remnants of an offline, analog time. Until the 2010s many households were still served through electromechanical Ferraris-style meters that have their origin in the late 19th century[8, 48, 21]. Today, under the umbrella term *Smart Metering*, the shift towards fully computerized, often networked meters is well underway. The roll out of these *Smart Meters* has not been very smooth overall with some countries severely lagging behind. As a safety-critical technology, smart metering technology is usually standardized on a per-country basis. This leads to an inhomogenous landscape with–in some instances–wildly incompatible systems. Often vendors only serve a single country or have separate models of a meter for each country. This complex standardization landscape and market situation has led to a proliferation of highly complex, custom-coded microcontroller firmware. The complexity and scale of this–often network-connected–firmware makes for a ripe substrate for bugs to surface.

A remotely exploitable flaw inside the firmware of a component of a smart metering system could have consequences ranging from impaired billing functionality to an existential threat to grid stability[1, 2]. In a country where meters commonly include disconnect switches for purposes such as prepaid tariffs, a coordinated attack could at worst cause widespread activation of grid safety systems through oscillations caused by repeated cycling of megawatts of load capacity at just the wrong frequency[50].

Mitigation of these attacks through firmware security measures is unlikely to yield satisfactory results. The enormous complexity of smart meter firmware makes firmware security extremely labor-intensive. The diverse standardization landscape makes a coordinated, comprehensive response unlikely.

In this paper, instead of focusing on the very hard task of improving firmware security we introduce a pragmatic solution to the–in our opinion likely–scenario of a large-scale compromise of smart meter firmware. In our concept the components of the smart meter that are threatened by remote compromise are equipped with a physically separate *safety reset controller* that listens for a "reset" command transmitted through the electrical grid's frequency and on reception forcibly resets the smart meter's entire firmware to a known-good state. Our safety reset controller receives commands through Direct Sequence Spread Spectrum (DSSS) modulation carried out on grid frequency through a large controllable load such as an aluminium smelter. After forward error correction and cryptographic verification it re-flashes the meter's main microcontroller over the standard JTAG interface. Note that our modulation technique is one *changing grid frequency itself*. This is fundamentally different in both generation and detection from systems such as traditional PLC that superimpose a signal on grid voltage, but leave the underlying grid frequency itself unaffected.

Starting from a high level architecture, we have carried out simulations of our concept's performance under real-world conditions. Based on these simulations we implemented an end-to-end prototype of our proposed safety reset controller as part of a realistic smart meter demonstrator. Finally, we experimentally validated our results and we will conclude with an outline of further steps towards a practical implementation.

This work contains the following contributions:

1. We introduce Grid Frequency Modulation (GFM) as a communication primitive.
2. We elaborate the fundamental physics underlying GFM and theorize on the constrains of a practical implementation.
3. We design a communication system based on GFM.
4. We carry out extensive simulations of our systems to determine its performance characteristics.

2 Related work

2.1 Security and Privacy in the Smart Grid

The smart grid in practice is nothing more or less than an aggregation of embedded control and measurement devices that are part of a large control system. This implies that all the same security concerns that apply to embedded systems in general also apply to the components of a smart grid. Where programmers have been struggling for decades now with issues such as input validation[34], the same potential issue raises security concerns in smart grid scenarios as well[36, 33]. Only, in smart grid we have two complicating factors present: Many components are embedded systems, and as such inherently hard to update. Also, the smart grid and its control algorithms act as a large partially distributed system making problems such as input validation or authentication harder[6] and adding a host of distributed systems problems on top[30].

Given that the electrical grid is essential infrastructure in our modern civilization, these problems amount to significant issues. Attacks on the electrical grid may have grave consequences[1, 33] while the long replacement cycles of various components make the system slow to adapt. Thus, components for the smart grid need to be built to a much higher standard of security than most consumer devices to ensure they live up to well-funded attackers even decades down the road. This requirement intensifies the challenges of embedded security and distributed systems security among others that are inherent in any modern complex technological system. The safety-critical nature of the modern smart metering ecosystem in particular was quickly recognized[1].

A point we will not consider in much depth in this work is theft of electricity. While in publications aimed towards the general public the introduction of smart metering is always motivated with potential cost savings and ecological benefits, in industry-internal publications the reduction of electricity theft is often cited as an incentive[11]. Likewise, academic publications tend to either focus on other benefits such as generation efficiency gains through better forecasting or rationalize the consumer-unfriendly aspects of smart metering with social benefits[40]. They do not usually point out revenue protection mechanisms as incentives[1, 2].

A serious issue in smart metering setups is customer privacy. Even though the meter "only" collects aggregate energy consumption of a whole household, this data is highly sensitive[37]. This counterintuitive fact was initially overlooked in smart meter deployments leading to outrage, delays and reduced features[10]. The root cause of this problem is that given sufficient timing resolution these aggregate measurements contain ample entropy. Through disaggregation algorithms, individual loads can be identified and through pattern matching even complex usage patterns can be discerned with alarming accuracy[23] in the same way that similar privacy issues arise in many other areas of modern life through other kinds of pervasive tracking and surveillance[52].

Another fundamental challenge in smart grid implementations is the central role of smart electricity meters in the smart grid ecosystem. Smart meters are

used both for highly-granular load measurement and in some countries also for load switching[51]. Smart electricity meters are effectively consumer devices. They are built down to a certain price point that is measured by the burden it puts on consumers and that is divided by the relatively small market served by a single smart meter implementation. Such cost requirements can preclude security features such as the use of a standard hardened software environment on a high powered embedded system. Landis+Gyr, a large manufacturer that makes most of its revenue from utility meters in their 2019 annual report write that they 36 % of their total R&D budget on embedded software while spending only 24 % on hardware R&D[31, 32], indicating a significant tension between firmware security and a smart meter vendor's bottom line.

2.2 The state of the art in embedded security

Embedded software security generally is much harder than security of higher-level systems. The primary two factors affecting this are that on one hand, embedded devices usually run highly customized firmware that (often by necessity) is rarely updated. On the other hand, embedded devices often lack advanced security mechanisms such as memory management units that are found in most higher-power devices. Even well-funded companies continue to have trouble securing their embedded systems. A spectacular example of this difficulty is the 2019 flaw in Apple's iPhone SoC first-stage ROM bootloader that allows for the full compromise of any iPhone before the iPhone X given physical access to the device[4]. iPhone 8, one of the affected models, was still being manufactured and sold by Apple until April 2020. In another instance in 2016, researchers found multiple flaws in the secure world firmware used by Samsung in their mobile phone SoCs. The flaws they found were both severe architectural flaws such as secret user input being passed through untrusted userspace processes without any protection as well as shocking cryptographic flaws such as CVE-2016-1919[4][26]. And Samsung is not the only large multinational corporation having trouble securing their secure world firmware implementation. In 2014 researchers found an embarrassing integer overflow flaw in the low-level code handling untrusted input in Qualcomm's QSEE firmware[43]. For an overview of ARM TrustZone including a survey of academic work and past security vulnerabilities of TrustZone-based firmware see [41].

If even companies with R&D budgets that rival some countries' national budgets at mass-market consumer devices have trouble securing their mass market secure embedded software stacks, what is a much smaller smart meter manufacturer to do? Especially if national standards mandate complex protocols such as TLS that are tricky to implement correctly[20], this manufacturer will be short on options to secure their product.

[4] `http://cve.circl.lu/cve/CVE-2016-1919`

2.3 Attack surface in the smart grid

From the incidents we outlined in the previous paragraphs we conclude that in smart metering technology, market incentives do not currently provide the conditions for a level of device security that will reliably last the coming decades. Considering this tension, in this paragraph we examine the cyberphysical risks that arise from attacks on the smart grid in the first place. These risks arise at three different infrastructure levels.

The first level is that of attacks on centralized control systems. This type of attack is often cited in popular discourse and to our knowledge is the only type of attack against an electric grid that has ever been carried out in practice at scale[33]. Despite their severity, these attacks do not pose a strictly *scientific* challenge since they are generic to any industrial control system. Their causes and countermeasures are generally well-understood and the hardest challenge in their prevention is likely to be budgetary constraints.

Beyond the centralized control systems, the next target for an attacker may be the communication links between those control systems and other smart grid components. While in some countries such as Italy special-purpose systems such as PLC are common[44], overall, IP-based technologies have proliferated according to the larger trend in commputing towards IP-based communications. This proliferation of IP-based communication links brings along the possibility for the application of generic network security measures from the IP world to the smart grid domain. In this way, a standardized, IP-based protocol stack unlocks decades of network security improvements at little cost.

Beyond these layers towards the core of the smart grid's control infrastructure, an attacker might also corrupt the network from the edges and target the endpoint devices itself. The large scale deployment of networked smart meters creates an environment that is favorable to such attacks.

2.4 Cyberphysical threats in the smart grid

Assuming that an attacker has compromised devices on any of these levels of smart grid infrastructure, what could they do with their newly gained power? The obvious action would be to switch off everything. Of all scenarios, this is both the most likely in practice—it is exactly what happened in the russian cyberattacks on the Ukranian grid[33]—but it is also the easiest to mitigate since the vulnerable components are few and centralized. Mitigations include the installation of fail-safes as well as a defense in depth approach to hardening the grid's cyber-infrastructure.

Another possible action for an attacker would be to forge energy measurements in an attempt to cause financial mayhem. Both individual consumers as well as the utility could be targeted by such an attack. While such an attack might have localized success, larger-scale discrepancies will likely quickly be caught by monitoring systems. For example, if a large number of meters in an area systematically under- or over-reported their energy readings, meter readings across the affected area would no longer add up with those of monitoring devices in other locations in the transmission and distribution grid.

In some countries, smart meter functionality goes beyond mere monitoring devices and also includes remotely controlled switches. There are two types of these switches: Switches to support *Demand-Side Management* (DMS) and cut off-switches that are used to punish defaulting customers. Demand Side Management is when a grid operator can remotely control the timing of large, non-time-critical loads on the customer's premises[14]. A typical example of this is a customer using an electric water heater: The heater is outfitted with a large hot water storage tank and is connected hooked up to the utility's DSM system. The customer does not care when exactly their water is heated as long as there is enough of it, and the utitliy offers them cheaper rates for the electricity used for heating in exchange for control over its precise timing. The utility uses this control to even out peaks in the consumption/production imbalance, remotely enabling DSM systems during off-peak times and disabling them during peak hours. In contrast to DSM, cut-off switches are switches placed in-between the grid and the entire customer's household such that the utility can disconnect non-paying customers without incurring the expense of sending a technician to the customer's premises. Unlike DSM systems, cut-off switches are not opt-in[1, 45]. An attack that uses cut-off switches would obviously immediately cause severe mayhem. Attacks on DSM may have more limited immediate impact as affected consumers may not notice an interruption for several hours.

Instead of switching off loads outright, an attack employing DSM switches (and potentially also cut-off switches) could choose to target the grid's stability. By synchronizing many compromised smart meters to switch on and off a large amount of load capacity, an attacker might cause the entire electrical grid to oscillate[28, 50, 27]. As a large system of coupled mechanical systems, the electrical grid exhibits a complex frequency-domain behavior. These resonance effects, colloquially called "modes", are well-studied in power system engineering[42, 22, 17, 9]. As they can cause issues even under normal operating conditions, a large effort is invested in dampening these resonances. Howewer, fully eliminating them under changing load conditions may not be achievable.

2.5 Communication Channels on the Grid

A core part of intervening with any such cyberattack is the ability to communicate remediary actions to the devices under attack. There is a number of well-established technologies for communication on or along power lines. We can distinguish three basic system categories: Systems using separate wires (such as DSL over landline telephone wiring), wireless radio systems (such as LTE) and *Power Line Communication* (PLC) systems that reuse the existing mains wiring and superimpose data transmissions onto the 50 Hz mains sine[24, 25].

During a large-scale cyberattack, availability of internet and cellular connectivity cannot be relied upon. An attacker may already have disabled such systems in a separate attack, or they may go down along with parts of the electrical grid. Traditional powerline communication systems or an utitly's proprietary wireless systems would work, but at a range of no more than several tens of kilometers

reaching all meters in a country would require a large upfront infrastructure investment.

3 Grid Frequency as a Communication Channel

We propose to approach the problem of broadcasting an emergency signal to all smart meters within a synchronous area by using grid frequency as a communication channel. Despite the awesome complexity of large power grids, the physics underlying their response to changes in load and generation is surprisingly simple. Individual machines (loads and generators) can be approximated by a small number of differential equations and the entire grid can be modelled by aggregating these approximations into a large system of nonu differential equations. As a consequence, small signal changes in generation/consumption power balance cause an approximately proportional change in frequency[29, 9, 47, 46]. This *Power Frequency Charactersistic* is about $25\,\mathrm{GW\,Hz^{-1}}$ for the continental European synchronous area according to European electricity grid authority ENTSO-E.

If we modulate the power consumption of a large load such as a multi-megawatt aluminium smelter, this modulation will result in a small change in frequency according to this characteristic. So long as we stay within the operational limits set by ENTSO-E[47, 18], this change will not degrade the operation of other parts of the grid. The advantages of grid frequency modulation are the fact that a single transmitter can cover an entire synchronous area as well as low receiver hardware complexity.

To the best of the authors' knowledge, grid frequency modulation has only ever been proposed as a communication channel at very small scales in microgrids before[49] and has not yet been considered for large-scale application.

3.1 Characterizing Grid Frequency

In utility SCADA systems, Phasor Measurement Units (PMUs, also called *synchrophasors*) are used to precisely measure grid frequency among other parameters. This task is much more complicated in practice than it might appear at first glance since a PMU has to make extremely precise measurements, track fast changes in frequency and handle even distorted input signals. Detail on the inner workings of commercial phasor measurement units is scarce but there is a large amount of academic research on sophisticated phasor measurement algorithms[38, 12, 5].

Since we do not need reference standard-grade accuracy for our application we chose to start with a very basic algorithm based on short-time fourier transform (STFT). Our system uses the universal frequency estimation approach of experimental physicists Gasior and Gonzalez at CERN[19]. The Gasior and Gonzalez algorithm[19] passes the windowed input signal through a DFT, then interpolates the signal's fundamental frequency by fitting a wavelet such as a Gaussian to the largest peak in the DFT results. The bias parameter of this curve fit is

an accurate estimation of the signal's fundamental frequency. This algorithm is similar to the simpler interpolated DFT algorithm used as a reference in much of the phasor measurement literature[7].

To collect ground truth measurements for our analysis of grid frequency as a communication channel, we developed a device to safely record real mains voltage waveforms. Our system consists of an `STM32F030F4P6` ARM Cortex M0 microcontroller that records mains voltage using its internal 12-bit ADC and transmits measured values through a galvanically isolated USB/serial bridge to a host computer. We derive our system's sampling clock from a crystal oven to avoid frequency measurement noise due to thermal drift of a regular crystal: 1 ppm of crystal drift would cause a grid frequency error of 50 µHz. We validated the performance of our crystal oven solution by benchmarking it against a GPS 1pps reference.

4 Grid Frequency Modulation

Given the grid characteristics we measured using our custom waveform recorder and a model of our transmitter, we can derive parameters for the modulation of our broadcast system. In its most basic form a transmitter for grid frequency modulation would be a very large controllable load connected to the power grid at a suitable vantage point. A spool of wire submerged in a body of cooling liquid such as a small lake along with a thyristor rectifier bank would likely suffice to perform this function during occasional cybersecurity incidents. We can however decrease hardware and maintenance investment even compared to this rather uncultivated solution by repurposing large industrial loads as transmitters. Going through a list of energy-intensive industries in Europe[15], we found that an aluminium smelter would be a good candidate. In aluminium smelting, aluminium is electrolytically extracted from alumina solution. High-voltage mains power is transformed, rectified and fed into about 100 series-connected electrolytic cells forming a *potline*. Inside these pots alumina is dissolved in molten cryolite electrolyte at about 1000 °C and electrolysis is performed using a current of tens or hundreds of Kiloampère. The resulting pure aluminium settles at the bottom of the cell and is tapped off for further processing.

Aluminium smelters are operated around the clock, and due to the high financial stakes their behavior under power outages has been carefully characterized by the industry. Power outages of tens of minutes up to two hours reportedly do not cause problems in aluminium potlines[16, 39]. Recently, even techniques for intentional power modulation without affecting cell lifetime or product quality have been devloped to take advantage of variable energy prices.[13, 16]. An aluminium plant's power supply is controlled to constantly keep all smelter cells under optimal operating conditions. Modern power supply systems employ large banks of diodes or SCRs to rectify low-voltage AC to DC to be fed into the potline[3]. Potline voltage is controlled through a combination of a tap changer and a transductor. Individual cell voltages are controlled by changing the physical distance between anode and cathode distance. In this setup, power can be

modulated fully electronically. Since this system does not have any mechanical inertia, high modulation rates can reasonably be achieved.

4.1 Parametrizing Modulation for GFM

Modulating 25 MW of smelter power would yield a frequency shift of 1 mHz. At an RMS frequency noise of around 10 mHz in the band around 1 Hz, this results in challenging SNR. Under such conditions, the obvious choice for modulation are spread-spectrum techniques. Thus, we approached the setting using Direct Sequence Spread Spectrum for its simple implementation and good overall performance. DSSS chip timing should be as fast as the transmitter's physics allow to exploit the low-noise region between 0.2 Hz to 2.0 Hz in the frequency noise spectrum while avoiding any of the grid's oscillation modes. Going past ≈ 2 Hz would put strain on the receiver's frequency measurement subsystem[5]. Using a spread-spectrum technique allows us to reduce the effect of interference by spurious tones. In addition, spreading our signal's energy over frequency also reduces the likelihood that we cause the grid to oscillate along any of its modes.

To test our proposed approach, we wrote a proof-of-concept modulator and demodulator in Python and tested this proof-of-concept prototype with data captured from our grid frequency sensor. Our simulations covered a range of parameters in modulation amplitude, DSSS sequence bit depth, chip duration and detection threshold. Figure 1 shows symbol error rate (SER) as a function of modulation amplitude with Gold sequences of several bit depths. As can be seen, realistic modulation amplitudes are in the range around 1 mHz. In the continental European synchronous area, this corresponds to a modulation power of approximately 25 MW. Figure 2 shows SER against detection threshold relative to background noise. Figure 3 shows SER against chip duration for a given fixed symbol length. As expected from looking at our measured grid frequency noise spectrum, performance is best for short chip durations and worsens for longer chip durations since shorter chip durations move our signals' bandwidth into the lower-noise region from 0.2 Hz to 2 Hz.

4.2 Parametrizing a proof-of-concept "Safety Reset" System Based on GFM

Taking these modulation parameters as a starting point, we proceeded to create a proof-of-concept smart meter emergency reset system. On top of the modulation described in the previous paragraphs we layered simple Reed-Solomon error correction[35] and some cryptography. The goal of our PoC cryptographic implementation was to allow the sender of an emergency reset broadcast to authorize a reset command to all listening smart meters. An additional constraint of our setting is that due to the extremely slow communication channel all messages should be kept as short as possible. The solution we chose for our PoC is a simplistic hash chain using the approach from the Lamport and Winternitz One-time Signature (OTS) schemes. Informally, the private key is a random bitstring. The public key is generated by recursively applying a hash function to

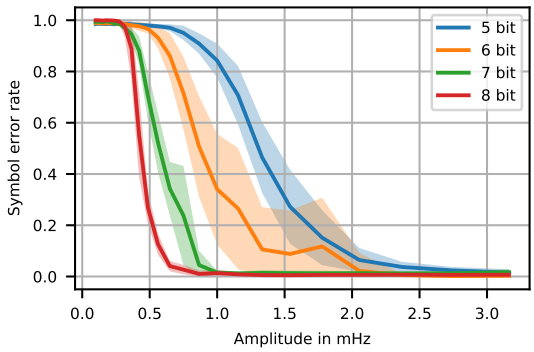

Fig. 1: Symbol Error Rate as a function of modulation amplitude for Gold sequences of several lengths.

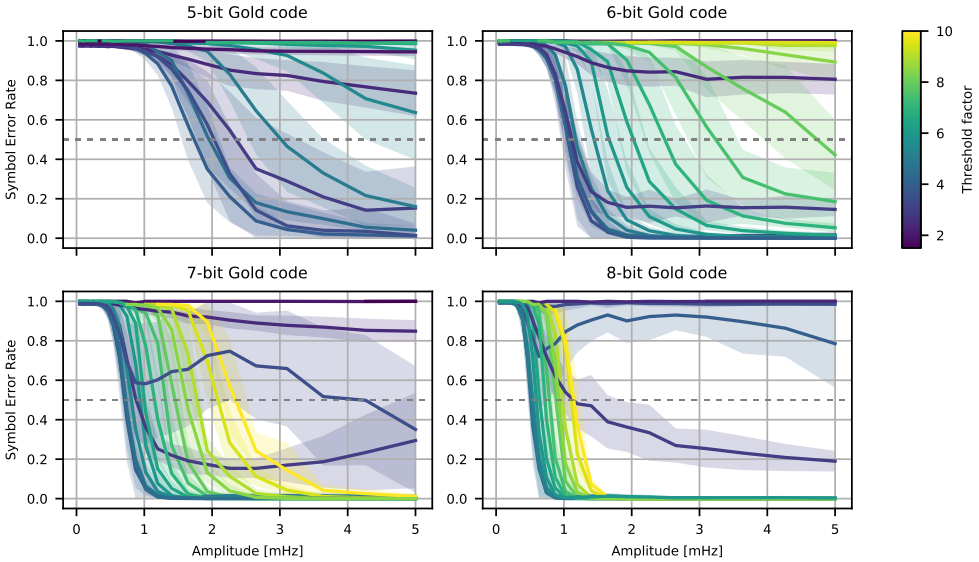

Fig. 2: SER vs. Amplitude and detection threshold. Detection threshold is set as a factor of background noise level.

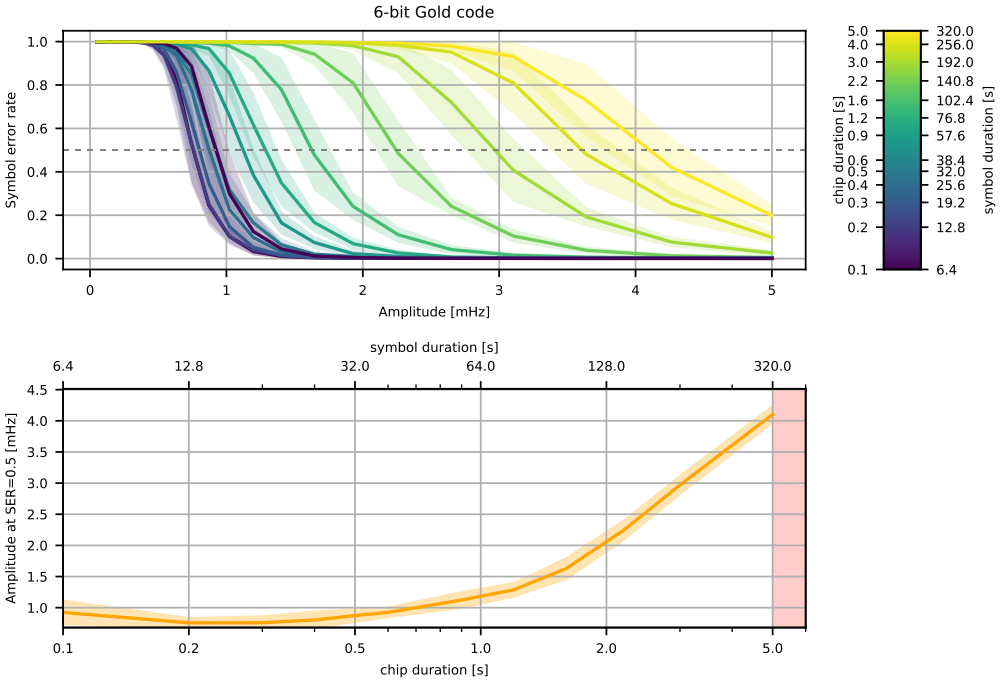

Fig. 3: SER vs. DSSS chip duration.

this key a number of times. Each smart meter reset command is then authorized by disclosing subsequent elements of this series. Unwinding the hash chain from the public key at the end of the chain towards the private key at its beginning, at each step a receiver can validate the current command by checking that it corresponds to the previously unknown input of the current step of the hash chain. Replay attacks are prevented by recording the most recent valid command. This simple scheme does not afford much functionality but it results in very short messages and removes the need for computationally public key cryptography inside the smart meter.

4.3 Experimental results

For a realistic proof of concept, we decided to implement our signal processing chain from DSSS demodulator through error correction up to our simple cryptography layer in microcontroller firmware and demonstrate this firmware on actual smart meter hardware, shown in Figure 4. In our proof of concept a safety reset controller is connected to the main application microcontroller of a smart meter. The reset controller is tasked with listening for authenticated reset

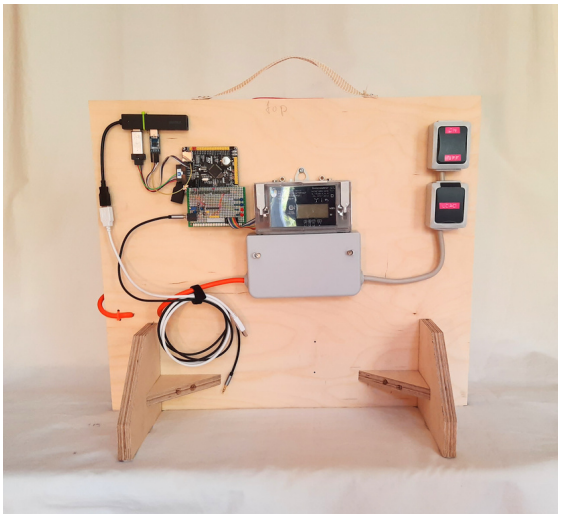

Fig. 4: The completed prototype setup. The board on the left is the safety reset microcontroller. It is connected to the smart meter in the middle through an adapter board. The top left contains a USB hub with debug interfaces to the reset microcontroller. The cables on the bottom left are the debug USB cable and the 3.5 mm audio cable for the simulated mains voltage input.

commands on the voltage waveform, and on reception of such a command resetting the smart meter application controller by flashing a known-good firmware image to its memory.

The signal processing chain of our PoC is shown in Figure 5. To interoperate with existing implementations of SHA-512 and reed-solomon decoding, this implementation was written in the C programming language. To demonstrate an application close to a field implementation, we chose an Easymeter `Q3DA1002` smart meter as our reset target. This model is popular in the German market and readily available second-hand. The meter consists of three isolated metering ASICs connected to a data logging and display PCB through infrared optical links. To demonstrate the safety reset's firmware reset functionality, we connected our safety reset microcontroller to the Texas Instruments `MSP430` microcontroller on the meter's display and data logging board through the JTAG debug interface that the board's vendor had conveniently left accessible. We ported part of `mspdebug`[5] to drive the meter microcontroller's JTAG interface and wrote a piece of demonstrator code that overwrites the meter's firmware with one that displays an identifying string on the meter's display after boot-up.

Since we did not have an aluminium smelter ready, we decided to feed our proof-of-concept reset controller with an emulated grid voltage sine wave from a computer's headphone jack. Where in a real application this microcontroller

[5] `https://dlbeer.co.nz/mspdebug/`

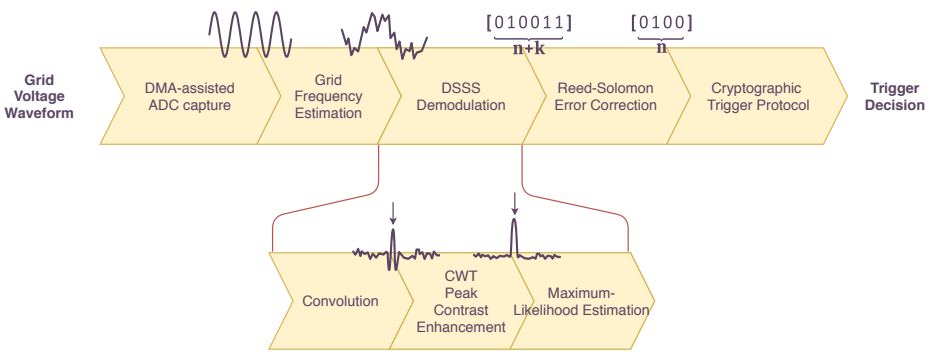

Fig. 5: The signal processing chain of our demonstrator.

might take ADC readings of input mains voltage divided down by a long resistive divider chain, we instead feed the ADC from a 3.5 mm audio input. For operational safety, we disconnected the meter microcontroller from its grid-referenced capacitive dropper power supply and connected it to our reset controlller's debug USB power supply.

We performed several successful experiments using a signature truncated at 120 bit and a 5 bit DSSS sequence. Taking the sign bit into account, the length of the encoded signature is 20 DSSS symbols. On top of this we used Reed-Solomon error correction at a 2:1 ratio inflating total message length to 30 DSSS symbols. At the 1 s chip rate we used in other simulations as well this equates to an overall transmission duration of approximately 15 min. To give the demodulator some time to settle and to produce more realistic conditions of signal reception we padded the modulated signal unmodulated noise on both ends.

5 Discussion

For our proof of concept, before settling on the commercial smart meter we first tried to use an `EVM430-F6779` smart meter evaluation kit made by Texas Instruments. This evaluation kit did not turn out well for two main reasons. One, it shipped with half the case missing and no cover for the terminal blocks. Because of this some work was required to get it electrically safe. Even after mounting it in an electrically safe manner the safety reset controller prototype would also have to be galvanically isolated to not pose an electrical safety risk since the main MCU is not isolated from the grid and the JTAG port is also galvanically coupled. The second issue we ran into was that the development board is based around a specific microcontroller from TI's `MSP430` series that is incompatible with common JTAG programmers.

Our initial assumption that a development kit would be easier to program than a commercial meter did not prove to be true. Contrary to our expectations the commercial meter had JTAG enabled allowing us to easily read out its stock firmware without either reverse-engineering vendor firmware update files nor

circumventing code protection measures. The fact that its firmware was only available in its compiled binary form was not much of a hindrance as it proved not to be too complex and all we wanted to know we found out with just a few hours of digging in Ghidra[6].

In the firmware development phase our approach of testing every module individually (e.g. DSSS demodulator, Reed-Solomon decoder, grid frequency estimation) proved to be very useful. In particular debugging benefited greatly from being able to run several thousand tests within seconds. In case of our DSSS demodulator, this modular testing and simulation architecture allowed us to simulate thousands of runs of our implementation on test data and directly compare it to our Jupyter/Python prototype. Since we spent more time polishing our embedded C implementation it turned out to perform better than our Python prototype while still exhibiting the same fundamental response to changes to its parameters. One significant bug we fixed in the embedded C version was the Python version's tendency towards incorrect decodings at even very large amplitudes.

In accordance with our initial estimations we did not run into any code space nor computation bottlenecks for chosing floating point emulation instead of porting over our algorithms to fixed point calculations. The extremely slow sampling rate of our systems makes even heavyweight processing such as FFT or our brute force dynamic programming approach to DSSS demodulation possible well within our performance constraints.

Since we are only building a prototype we did not optimize firmware code size. At around 64 kB, the compiled code size of our firmware implementation is slightly larger than we would like. The overall most heavy-weight operations are the SHA512 implementation from libsodium and the FFT from ARM's CMSIS signal processing library. Especially the SHA512 implementation has large potential for size optimization because it is highly optimized for speed using extensive manual loop unrolling. Despite being larger than what we initially targeted, this firmware is still small compared to the firmware space available in commercially deployed smart meters. We estimate that even without additional optimizations, our PoC firmware is already within the realm of firmware size that could be implemented in a commercially viable safety reset controller.

6 Conclusion

In this paper we have developed an end-to-end design of a reset system to restore smart meters to a safe operating state during an ongoing large-scale cyberattack. To allow our system to be triggered even in the middle of a cyberattack we have developed a broadcast data transmission system based on intentional modulation of global grid frequency. We have shown the viability of our end-to-end design through simulations. To put these simulations on a solid foundation we have developed a grid frequency measurement methodology comprising of a custom-designed hardware device for electrically safe data capture and a set of software

[6] `https://ghidra-sre.org/`

tools to archive and process captured data. Our simulations show good behavior of our broadcast communication system and give an indication that cooperating with a large consumer such as an aluminium smelter would be a feasible way to set up a transmitter with low hardware overhead. We have outlined a simple cryptographic protocol ready for embedded implementation in resource-constrained systems that allows triggering a safety reset with a response time of less than 30 minutes. We have experimentally validated our system using simulated grid frequency data in a demonstrator setup based on a commercial microcontroller as our safety reset controller and an off-the-shelf smart meter. Source code and electronics CAD designs are available at the public repository listed at the end of this document.

References

[1] Ross Anderson and Shailendra Fuloria. "Who controls the off switch?" In: *2010 First IEEE International Conference on Smart Grid Communications*. Gaithersburg, MD, 2010, pp. 96–101. DOI: 10.1109/SMARTGRID.2010.5622026.

[2] Ross J. Anderson. *Security engineering. A guide to building dependable distributed systems*. 3rd. Preview of upcoming edition. Wiley, 2020.

[3] Mohammed W. Ayoub and Francis V. P. Robinson. *A comparative study between diode and thyristor based AC to DC converters for aluminium smelting process*. 2013. DOI: https://doi.org/10.1109/IEEEGCC.2013.6705851.

[4] Leo Becker. *checkm8: Boot-Exploit soll neuere iPhones knacken*. Ed. by Heise Online. Sept. 27, 2019. URL: https://www.heise.de/mac-and-i/meldung/checkm8-Boot-Exploit-soll-neuere-iPhones-knacken-4542075.html.

[5] Daniel Belega and Dario Petri. "Accuracy Analysis of the Multicycle Synchrophasor Estimator Provided by the Interpolated DFT Algorithm". In: *IEEE Transactions on Instrumentation and Measurement* 62 (5 2013), pp. 942–953. ISSN: 0018-9456. DOI: 10.1109/tim.2012.2236777.

[6] Matt Blaze et al. "The role of trust management in distributed systems security". In: *Secure Internet Programming*. Springer, 1999, pp. 185–210.

[7] Jozef Borkowski, Dariusz Kania, and Janusz Mroczka. "Interpolated-DFT-Based Fast and Accurate Frequency Estimation for the Control of Power". In: *IEEE Transactions on Industrial Electronics* 61 (12 2014), pp. 7026–7034. ISSN: 0278-0046. DOI: 10.1109/tie.2014.2316225.

[8] Stuart Borlase, ed. *Smart Grids: Advanced Technologies and Solutions*. Electric Power and Energy Engineering. CRC Press, 2017. ISBN: 978-1-4987-9955-3.

[9] Valentin Crastan. *Elektrische Energieversorgung 3*. 2012. ISBN: 978-3-642-20099-1. DOI: 10.1007/978-3-642-20100-4.

[10] Colette Cuijpers and Bert-Jaap Koops. "Smart metering and privacy in Europe: lessons from the Dutch case". In: *European data protection. Coming of age* (2012), pp. 269–293. DOI: https://doi.org/10.1007/978-94-007-5170-5_12.

[11] R. Czechowski and A. M. Kosek. "The most frequent energy theft techniques and hazards in present power energy consumption". In: *2016 Joint Workshop on Cyber- Physical Security and Resilience in Smart Grids (CPSR-SG)*. IEEE. Apr. 2016, pp. 1–7. DOI: 10.1109/CPSRSG.2016.7684098.

[12] Asja Derviškadić, Paolo Romano, and Mario Paolone. "Iterative-interpolated DFT for synchrophasor estimation: A single algorithm for P-and M-class compliant PMUs". In: *IEEE Transactions on Instrumentation and Measurement* 67.3 (2017), pp. 547–558.

[13] Roman Düssel. "Paradigm Shift in the Indication of a Stable Cell during Power Modulation". In: *Proceedings of the 36th International ICSOBA Conference*. 2018.

[14] Dacfey Dzung, Inigo Berganza, and Alberto Sendin. "Evolution of powerline communications for smart distribution: From Ripple Control to OFDM". In: *2011 IEEE International Symposium on Power Line Communications and Its Applications* (2011). DOI: 10.1109/ISPLC.2011.5764444.

[15] Christian Egenhofer et al. *Composition and Drivers of Energy Prices and Costs: Case Studies in SelectedEnergy Intensive Industries – 2018*. 2018. DOI: 10.2873/937326.

[16] David Eisma and Pretesh Patel. "Challenges in Power Modulation". In: *Essential Readings in Light Metals, Volume 2, Aluminum Reduction Technology*. Ed. by Geoff Bearne, Marc Dupuis, and Gary Tarcy. 2016, pp. 683–688.

[17] ENTSO-E System Protection Dynamics and WG. *Oscillation Event 03.12.2017*. Mar. 2018.

[18] ENTSO-E Working Group Incident Classification Scale Under System Operations Committee. *Incidents Classification Methodology*. 2014.

[19] M Gasior and JL Gonzalez. *Improving FFT frequency measurement resolution by parabolic and gaussian interpolation*. 2004.

[20] Martin Georgiev et al. "The Most Dangerous Code in the World: Validating SSL Certificates in Non-Browser Software". In: *ACM Conference on Computer and Communications Security*. 2012, pp. 38–49.

[21] German Government Bundesnetzagentur. *Monitoring Report 2018*. 2018.

[22] E. Grebe et al. "Low Frequency Oscillations in the Interconnected System of Continental Europe". In: IEEE, Aug. 2010. DOI: 10.1109/PES.2010.5589932{\textperiodcentered}.

[23] Ulrich Greveler et al. "Multimedia Content Identification Through Smart-Meter Power Usage Profiles". In: *Computers, Privacy and Data Protection* (2012).

[24] Vehbi C. Güngör et al. "Smart Grid Technologies: Communication Technologies and Standards". In: *IEEE Transactions on Industrial Informatics* 7.4 (Nov. 2011), pp. 529–539.

[25] Yasin Kabalci. "A survey on smart metering and smart grid communication". In: *Renewable and Sustainable Energy Reviews* 57 (2016), pp. 302–318. DOI: 10.1016/j.rser.2015.12.114.

[26] Uri Kanonov and Avishai Wool. "Secure containers in Android: the Samsung KNOX case study". In: *Proceedings of the 6th Workshop on Security and Privacy in Smartphones and Mobile Devices*. 2016, pp. 3–12.

[27] Jinsub Kim and Lang Tong. "On Topology Attack of a Smart Grid: Undetectable Attacks and Countermeasures". In: *IEEE Journal on Selected Areas in Communications* 31.7 (July 2013). DOI: 10.1109/JSAC.2013.130712.

[28] Oliver Kosut et al. "Malicious Data Attacks on the Smart Grid". In: *IEEE Transactions on Smart Grid* 2.4 (Nov. 2011), pp. 645–658.

[29] Prabha Kundur. *Power system stability and control*. eng. The EPRI power system engineering series. New York, NY u.a.: McGraw-Hill, 1994. ISBN: 007035958X.

[30] Leslie Lamport, Robert Shostak, and Marshall Pease. "The Byzantine Generals Problem". In: *ACM Transactions on Programming Languages and Systems* 4.3 (July 1982), pp. 382–401.

[31] Landis+Gyr Group AG. *Landis+Gyr Annual Report 2019*. May 28, 2020.

[32] Landis+Gyr Group AG. *Landis+Gyr Financial Report 2019*. May 6, 2020.

[33] Robert M. Lee, Michael J. Assante, and Tim Conway. "Analysis of the cyber attack on the Ukrainian power grid". In: *Electricity Information Sharing and Analysis Center (E-ISAC)* (2016).

[34] Nancy G. Leveson and Clark S. Turner. "An Investigation of the Therac-25 Accidents". In: *IEEE Computer* 26.7 (July 1993), pp. 18–41.

[35] David J. C. MacKay. *Information theory, inference, and learning algorithms*. Repr. with corr. Literaturverz. S. 613 - 619. Cambridge [u.a.]: Univ. Press, 2005. XII, 628. ISBN: 0521642981.

[36] Yilin Mo et al. "Cyber-Physical Security of a Smart Grid Infrastructure". In: *Proceedings of the IEEE* 100.1 (Jan. 2012), pp. 195–209.

[37] Andrés Molina-Markham et al. "Private Memoirs of a Smart Meter". In: *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building* (2010). Title from The ACM Digital Library.

[38] Claudio Narduzzi et al. "Fast-TFM—Multifrequency phasor measurement for distribution networks". In: *IEEE Transactions on Instrumentation and Measurement* 67.8 (2018), pp. 1825–1835.

[39] Harald A. Øye. *Power Failure, Temporary Pot Shut-Down, Restart and Repair*. 2012.

[40] McDaniel Patrick and McLaughlin Stephen. "Security and Privacy Challenges in the Smart Grid". In: *Secure Systems* (May 2009).

[41] Sandro Pinto and Nuno Santos. "Demystifying Arm TrustZone: A Comprehensive Survey". In: *ACM Comput. Surv.* 51.6 (Jan. 2019). ISSN: 0360-0300. DOI: 10.1145/3291047.

[42] Graham Rogers. "Power System Oscillations". In: Kluwer, 2000.

[43] Dan Rosenberg. "Qsee trustzone kernel integer over flow vulnerability". In: *Black Hat conference*. 2014.

[44] *Single Market Progress Report: Country Profiles – Italy*. Research rep. 2014.

[45] William G. Temple, Binbin Chen, and Nils Ole Tippenhauer. "Delay Makes a Difference: Smart Grid Resilience Under Remote Meter Disconnect Attack". In: *2013 IEEE International Conference on Smart Grid Communications*. 2013. DOI: https://doi.org/10.1109/SmartGridComm.2013.6688001.

[46] UCTE/ENTSO-E. *Operation Handbook*. 2004.

[47] UCTE/ENTSO-E. *Operation Handbook*. 2009.

[48] UK Department for Business Energy and Industrial Strategy. *Smart Meter Statistics Quarterly Report to end March 2019*. 2019.

[49] Andoni Urtasun et al. "Energy management strategy for a battery-diesel stand-alone system with distributed PV generation based on grid frequency modulation". In: *Renewable Energy* 66 (Jan. 2014), pp. 325–336.

[50] Yongdong Wu et al. "Resonance Attacks on Load Frequency Control of Smart Grids". In: *IEEE Transactions on Smart Grid* 9.5 (Sept. 2018), pp. 4490–4502. DOI: 10.1109/TSG.2017.2661307.

[51] Jixuan Zheng, David Wenzhong Gao, and Li Lin. "Smart meters in smart grid: An overview". In: *2013 IEEE Green Technologies Conference (GreenTech)*. IEEE. 2013, pp. 57–64.

[52] Shoshana Zuboff. *The Age of Surveillance Capitalism*. 2019.

This is version v0.7-25-g31cc8fa-dirty of this paper, generated on April 23, 2021. The git repository can be found at:

https://git.jaseg.de/safety-reset.git