International Journal of Information Security manuscript No.
(will be inserted by the editor)

Side-channel attacks and countermeasures for heart rate retrieval from ECG characterization device

Pablo Perez-Tirador · Madhav Desai · Alejandro Rodriguez · Elena Berral · Teresa Romero · Gabriel Caffarena · Ruzica Jevtic

Received: date / Accepted: date

Abstract With a rapid advance of the technology, side-channel attacks are gaining more attention in the security evaluation of electronic devices. The impact of the attacks on medical devices can be very dangerous: from retrieving private health data to attacking implantable devices causing life-threatening situations. This paper investigates the vulnerabilities of ECG characterization devices to power side-channel attacks and proposes power supply voltage modulation as a countermeasure. Experimental results indicate that random voltage modulation can effectively obscure heart rate retrieval from leaked power signals. Sine modulation is less effective and can be canceled by demodulation at the sine frequency visible in the leaked signal spectrum.

Keywords ECG · power side-channel attacks · voltage modulation · Hermite characterization

Pablo Perez-Tirador
Escuela Politecnica Superior, Universidad San Pablo- CEU, CEU Universities, Urbanización Montepríncipe s/n, 28668 Alcorcon, Madrid, Spain
E-mail: pablo.pereztirador@ceu.es

Madhav Desai
Department of Electrical Engineering IIT Bombay, Powai Mumbai 400 076, India

Alejandro Rodriguez
Escuela Politecnica Superior, Universidad San Pablo- CEU, CEU Universities, Urbanización Montepríncipe s/n, 28668 Alcorcon, Madrid, Spain

Elena Berral
Escuela Politecnica Superior, Universidad San Pablo- CEU, CEU Universities, Urbanización Montepríncipe s/n, 28668 Alcorcon, Madrid, Spain

Teresa Romero
Escuela Politecnica Superior, Universidad San Pablo- CEU, CEU Universities, Urbanización Montepríncipe s/n, 28668 Alcorcon, Madrid, Spain

Gabriel Caffarena
Escuela Politecnica Superior, Universidad San Pablo- CEU, CEU Universities, Urbanización Montepríncipe s/n, 28668 Alcorcon, Madrid, Spain

Ruzica Jevtic
Escuela Técnica Superior de Ingenieros de Telecomunicación, Universidad Politécnica de Madrid, Av. Complutense 30, 28040 Madrid, Spain

Author contributions

P.P-T. and R.J. wrote the main manuscript text, prepared the figures, supervised the lab work and performed analysis and interpretation of results. M.D. and G.C. developed the FPGA application tested in the paper and provided feedback on its working to understand the data. P.P-T. modified the application firmware to allow for modulation. A.R., E.B. and T.R. performed lab measurements, and processed the data in different scenarios. All authors reviewed the manuscript.

1 Introduction

Health data is considered sensitive information subject to strict protection policies, so it has become an attractive target for hackers. Healthcare data breaches have the highest cost of all industries and since 2020 these costs have increased by 53.3% [1]. Remote patient monitoring continues to evolve and rely on a use of small electronic devices to collect and process medical data. The devices are small since they are wearable and are unable to fit highly secure cryptographic algorithms. They are also more physically accessible than servers and more complex computers. This makes them particularly vulnerable to side-channel attacks that retrieve private data from electronic device physical leaks, such as power, electromagnetic radiations, sound and time.

Electromagnetic side-channel attacks are performed by holding an antenna in the vicinity of the device. This allows the attacker to capture electromagnetic radiation from

it and identify certain patterns in the leaked signal that reveal sensitive information processed in the device. The work in [2] demonstrates that a password that is typed into a mobile device can be retrieved even from longer distances (of a couple of meters) and a wall separation by infiltrating a small program inside the device and measuring its EM radiations afterwards. Low-power microcontrollers implemented together with radio transceivers on the same chip to enable wireless communications such as Bluetooth and WiFi, have been attacked from distances of up to 15 meters [3]. The attack was based on the coupling of data leakages from the digital to the analog/RF part of the chip, where they are amplified, modulated and transmitted along with legitimate data.

Power attacks usually require a small resistor to be inserted in the ground line of the device. By measuring the voltage over the resistor, the attacker can obtain sensitive information. These attacks have been performed on smart cards [4,5], laptops [6], and numerous implementations of AES cryptograhic algorithm [7–10]. Works on how to detect if the device is under attack have been reported in [11, 12].

Despite the catastrophic consequences that side-channel attacks could have on wearable and implantable devices, the work on this topic is very limited [13]. Information on ECG has been targeted in [14], attacks on pacemaker and heart defibrillator device were recently reported in [15–17] and acoustic attacks on communication between a smartphone and medical device in [18].

It is clear that heart devices are among the main targets of the attacks, as electrocardiogram (ECG) is one of the key physiological signals measured by the wearables. Widespread use of wearable ECG devices is needed to help reduce cardiovascular disease, which is the leading cause of death worldwide today [19]. Automatic detection of heart beats is critical to support ECG analysis, a tedious task usually performed as a visual inspection by a physician. As the implementation of these applications focuses more on portable devices for real-time processing, they become targets for the aforementioned attacks.

In this work we design a power side-channel attack on an application for automatic ECG signal characterization [20] implemented on a CW305 Field Programmable Gate Array (FPGA) platform. The characterization is based on special mathematical functions named Hermite polynomials. The ECG samples are fitted with a linear combination of these functions. The coefficients that multiply the functions are used for heart beat representation.

By analyzing the power during the heart beat characterization, we obtain characteristics of the patients' heart rate. We apply two different equipment settings for the attacks to anticipate dangerous scenarios. We then implement voltage modulation to the device's power supply to hide its activity and prevent the attacker from obtaining sensitive data from the ECG. Two different modulations are applied: random and sine, which are enabled by programming the microcontroller in charge of voltage regulation on the board. To the best of our knowledge, this is the first work to design and implement power attacks on ECG real-time monitoring devices.

2 Methodology

In this section we first describe the algorithm that is used for heartbeat characterization as we need to analyze which fraction of the algorithm is likely to cause the most leakage that can be exploited in the attack. We then give a more detailed explanation of the power side-channel attacks and the methodology that has been used in this work to retrieve sensitive information. Finally, we briefly describe countermeasures based on voltage modulation.

2.1 Heart rate characterization

The modelling and classification of heartbeats is a very important step in the processing of the electrocardiogram (ECG) [20–25]. However, it is a computationally costly and a time-consuming operation due to the vast amount of data needed [20,21,24]. Heartbeats, as registered by the ECG, show a distinct waveform consisting of 5 key points, named with letters, as shown in Fig. 1. The main idea behind the heartbeat characterization is to describe the heartbeat with a mathematical function and then look at the values of the function parameters to find out whether the heartbeat represents some anomaly. The patient diagnosis can thus be automated and help patients receive medical treatment on time. Hermite polynomials minimize the number of dimensions needed for ECG classification while preserving the beat morphology [20]. Unlike most ECG models, Hermite polynomials reflect the QRS shape, which carries important clinical information [23]. The centers of the heart waves (QRS complex) as well as their shape can be easily extracted and used for identifying different arrhythmia types. Additionally, Hermite polynomials were successfully used in ECG data compression [26], QRS complex clustering [27] and detection of miocardial infarction [28].

Mathematically speaking, Hermite polynomials are an orthogonal polynomial sequence that is used to approximate complex functions. Heartbeat representation with Hermite polynomials consists in reducing the mean square error (MSE) between the sampled QRS complex from a real ECG database and a polynomial approximation in successive iterations. Each approximation of the heartbeat is calculated optimizing the coefficients for a base polynomial and a time-scaling factor to adjust the width of the function to the actual beats. To

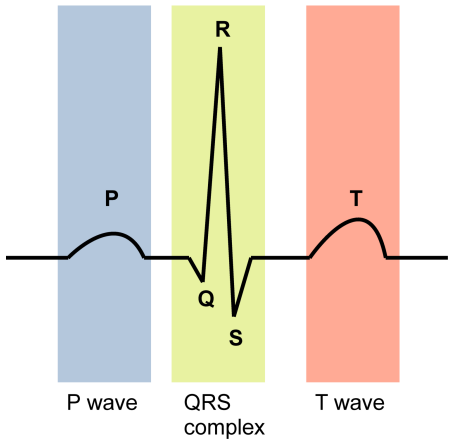

Fig. 1 Typical ECG for a single heartbeat.

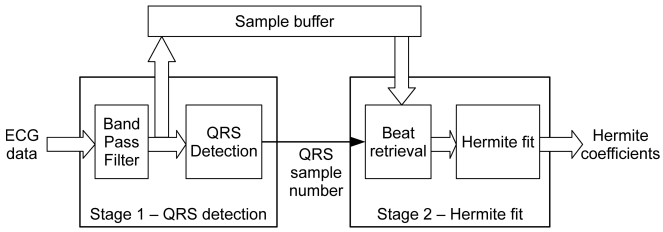

Fig. 2 Hermite application and coordination between stages.

achieve this, the window in which the heartbeats exist must be identified beforehand, calculating the center of the QRS complex.

We use an existing characterization application implemented in an FPGA, which allows for parallel execution of algorithms in hardware. The application consists of two stages, as illustrated in Fig. 2 [20]: filtering and peak detection (front-end) and computation of Hermite fits (back-end). The two stages communicate through a sample buffer. The front-end stage receives the heartbeat samples from the host at a constant rate and is always active, as it searches for the heartbeats. Its contribution to the power consumption is continuous, giving a baseline for power measurements. The first stage signals the arrival of a heartbeat to the back-end stage, which can only perform computations when a beat arrives. This stage reads samples from the buffer and, once all coefficients are calculated, sends the results to the host PC. The second stage operates in bursts creating peaks of activity in the power. This paves a way to possible side-channel attacks that aim to extract the original heart rate by detecting changes in the activity of the chip. Although we target a Hermite application, other real-time algorithms for ECG characterization [29–32] also create peaks of power activity that correspond to the heart rate and can be equally vulnerable to the attacks presented here.

2.2 Power side-channel attacks

Contrary to remote cyber attacks, side-channel attacks do not rely on mathematical properties of the cryptographic algorithm, but exploit physical leaks of the device instead. They rely on the fact that all sensitive data can be exposed through the capture of a device's sound emissions [33], execution time [34], power [35] or electro-magnetic radiations [10]. While there are a multitude of methods to hack medical devices, side-channel attacks can be the most cunning since even conventional cryptographic systems, the basis of most security measures, inherently leak side-channel information due to the fundamental nature of integrated circuits.

Side-channel attacks can be classified as cryptographic and non-cryptographic [36]. The former aim to retrieve the secret key of the cryptographic algorithm, while the latter search for correlation between the leaked signal and the activity inside the device to retrieve sensitive features of the data. Both types of attacks usually start with a simple power analysis (SPA) in the time and frequency domains to identify the key parts of the leaked signal that reveal critical information. Non-cryptographic attacks search for a correlation between the leaked signal and the activity inside the device. Cryptographic attacks are based on the collection of power traces for many different inputs and search for correlation between the power samples at the same instants in time in order to retrieve a secret key.

Side-channel attacks can be further classified as passive and active [36]. Passive attacks aim at retrieving sensitive information from unintentional leaked signals. The target for passive attacks could be any private health information such as: heart rate, heart rate variability, the QRS of the heartbeats or EEG frequency bands. For example, processor memory accesses are exploited to retrieve important information about the ECG in [14], while acoustic attacks on vibration-based channel between a smartphone and medical device are described in [18]. The analogy between pasive side-channel attacks and contactless medical monitoring is discussed in [37]. Active attacks aim to tamper with the device and affect its behavior, rendering it unavailable or making it behave abnormally and with fatal consequences. For example, in [15–17] it is shown that an attacker can alter the data processed inside a pacemaker or a heart defibrillator in the Bluetooth proximity of the device.

The attack presented here is passive and non-cryptographic, uses power consumption as the leaked signal and focuses on a hardware application, instead of specific registers or memory accesses. The power measurements are processed in order to retrieve a patient's heart rate. By knowing the heart rate has a low frequency, the following steps are taken:

1. The measured voltage is squared to obtain the signal power.

2. The resulting signal is filtered by a low pass filter since the heart rate is a signal that has low frequency.
3. The FFT is calculated from the filtered signal. The spectrum contains all frequency components of the signal among which the heart rate frequency is expected to be one of the dominant ones.
4. The resulting FFT is visually inspected to spot the heart rate.

Two different filters are applied in the second step: a Butterworth low pass filter and a median filter, to explore which filter gives better results. The preceding steps are applied in two ways: to the signal saved and recorded by the oscilloscope, processing it in Matlab, and by directly applying mathematical functions available in the oscilloscope.

2.3 Voltage modulation countermeasures

Voltage modulation was proven to be a promising countermeasure against both cryptographic and non-cryptographic side-channel attacks ([9,38,39,10,40]). By changing the voltage, the power signal available to the attacker is altered. In fact, it was seen that the voltage modulates the activity inside the device in such a way that if voltage variation is chosen correctly, its spectrum will overlap with the spectrum of the heart beat signal. In other words, we use voltage modulation to create aliasing of the load signal similar to the work in [9]. By changing the voltage, aliasing is created in the load spectrum to disable the attacker from retrieving private data and the heart rate frequency ceases to be one of the dominant ones, disabling the attacker from retrieving private data. Two different voltage waveforms are tested: random and sine. For random voltage, we use different time resolutions, while for the sine voltage we change the modulation frequency.

It should be noted that the attack presented here does not target the cryptographic stage of the data processing, but the computations done with plain text. Therefore, countermeasures such as noise injection or algorithmic masking are not applicable. One alternative measure that could be applied to mask the activity of the device is the introduction of noise generators, hardware blocks that can be implemented inside the FPGA and work in parallel to the main processing blocks and have been used also in cryptographic systems [41,42]. They perform meaningless activity in random bursts, producing extra peaks in the spectrum. While it does not need external components, like the power regulator, this technique occupies FPGA resources with blocks that do not perform real computation, and can affect the overall power consumption of the device. Voltage modulation, on the other hand, can be done with external components, saving computational resources. These two techniques are not incompatible and could be implemented side by side.

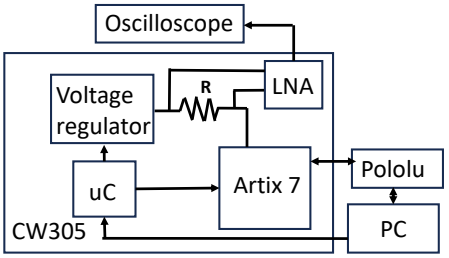

Fig. 3 Experimental setup.

3 Experimental setup

A block diagram of the experimental setup is shown in Fig. 3. This experimental setup was designed to perform the side-channel attacks and evaluate the effectiveness of different voltage modulation techniques in mitigating these attacks on ECG devices. We use a CW305 FPGA board as a victim device [43]. A hardware module with the Hermite application is implemented for an Artix 7 FPGA chip by using the AHIR tool chain [44]. The module uses Pololu UART-to-USB [45] to communicate with the PC from where the heartbeats are streamed from. For every heart beat, the fitting coefficients are extracted and sent back to the PC via UART.

3.1 ECG database

The ECG recordings are downloaded from the MIT-BIH arrhythmia database [46]. Two-channel ECG recordings are sampled at a frequency of 360 Hz and each recording has around 2000 heartbeats. We use 5 recordings from healthy and 5 from arrhythmia patients to make sure the attack is able to detect the differences in the heart rate between the two. Information about each recording, as well as classification of whether the recording captures arrhythmia episodes, is available at the database website. For example, synus arrhythmia was detected in recordings 108, 113, 115 and 123, while atrial fibrillation, an irregular and very rapid heart rhythm, was spotted in the recording 201 . Normal heartbeats were reported in recordings 100, 101, 116, 119 and 220. Some of these recordings had ventricular ectopy detected, but it was spotted in a very small number of heartbeats and the impact on the overall heart rate was insignificant.

3.2 Data transmission

Because the beats are pre-recorded, the samples are sent from the PC at a faster rate than they were recorded. The transmission time for each sample of one milisecond is configured on the PC side. However, one 16-bit sample requires two 8-bit UART packets, needing approximately 138 μs to be sent. Additionally, one full heartbeat uses around 250

samples (this number varies for normal and arrhythmia heartbeats), and the transmission protocol also includes dead time, which needs to be accounted for. As soon as one heartbeat is identified, the data reception does not stop, so this flow is continuous. With all this taken into account, a throughput of approximately 3 beats/s and 4 beats/s is observed for the healthy and arrhythmia patient recordings, respectively. For this mode of transmission, the processing rate is higher than the standard heart rate values (i.e. 1-1.7 beats/s).

The UART is configured at a 115200 bits per second baud rate. This baud rate is chosen to satisfy the heart rate transmission speed. It can be observed that it is higher than the rate that is needed for the transmission (i.e. 16 bits every milisecond equivalent to 16000 bits per second) and was configured in such a way to support faster transmission if one is needed. The transmission speed can be changed at the PC side by changing the transmission time between the samples.

3.3 Power measurements

A shunt resistor of 10 mΩ placed between the 1V voltage regulator and the FPGA chip is used for the attack. The voltage over the resistor is measured at the output of an on-board Low Noise Amplifier (LNA) to calculate the power.

For the online phase, power measurements are recorded by using an KeySight InfiniiVision MSOX3024T oscilloscope. They can be processed offline following the steps described in the previous section. We process the signal in two different ways since we noticed that only 64K samples are recorded when the signal is saved on a USB key attached to the oscilloscope. For low frequency signals such as a heart rate, a recording time of several seconds is needed to detect peaks in the activity and obtain significant information about signal features, like the heart rate in this case. However, by increasing the recording duration, the sample time resolution becomes poor due to the fixed number of samples. For this reason, we also process the signal directly by using mathematical functions available in the oscilloscope. The sampling rate is around 400 Ks/s, much higher than the sampling rate for the recorded signal.

3.4 Setup for countermeasures

For countermeasures we use CW305's on-board reconfigurable voltage regulator. The regulator is a Texas Instruments TPS56520PWPR [47] controlled by the microcontroller available on the board via an I^2C serial interface. With CW305's board configuration, the time latency for the voltage change is 1.6 ms, which limits voltage modulation to that sample rate.

CW305's microcontroller, a Microchip SAM3U2E [48], is responsible for connecting the target board with the host computer, programming the FPGA and configuring the on-board chips, including the PLLs and voltage regulator. Its firmware is open source and available from the manufacturer's repositories and natively offers an interface for sending commands to the board. This way, the host computer can set the FPGA's supply voltage and clock frequency via USB. The original firmware was modified to add a custom command that enables voltage modulation. After receiving this new command, a timer interrupt is set to update the FPGA's supply voltage at the maximum rate possible. The voltage values are read from a lookup table, which can contain any arbitrary waveform and sent to the voltage regulator in sequence. Using the voltage modulation command, the modulation frequency can be modified, allowing for divisions of the original frequency (with a minimum of 1.6 ms per table sample, as discussed earlier). For this paper we have tested two waveforms: a sine wave with 40 samples and a pseudorandom sequence with 5000 samples, both generated externally using Matlab (see Fig. 4). Voltage modulation is turned on and off by issuing commands from the host computer, so the FPGA can be programmed at a constant voltage and modulation can be used later, during measurements.

The maximum sine frequency is limited by the fastest change between the samples that the on-board microcontroller can handle. With 40 samples for the sine wave, and sample period of 1.6 ms between each sample, the maximum frequency for the sine wave is equal to $\frac{1}{40 \times 1.6 \text{ ms}} = 15.625$ Hz. We start our experiments with this frequency and gradually lower it down to 1 Hz.

For random wave, we also start with the minimum sample period of 1.6 ms between the samples of the random wave, and gradually increase this sample period to 10 s.

The lower limit for the sine wave, as well as the upper limit for the random wave sample period were established based on the experimental results as will be explained in detail in the next section. It was observed that further decreasing of the sine wave frequency and increasing the random wave sampling period was unable to mask the heart rate and therefore, was inefficient as a security countermeasure.

4 Experimental results

4.1 Attack results

There are two attack scenarios: recording the signal and processing it in Matlab (Mat scenario) and using the oscilloscope functions to process the signal directly (Osc scenario). The Hermite application calculates the polynomial fits as soon as it receives a heartbeat. The fits are sent back to the PC approximately at the heartbeat frequency. We use the UART's transmission line signal (TX on the FPGA side) to

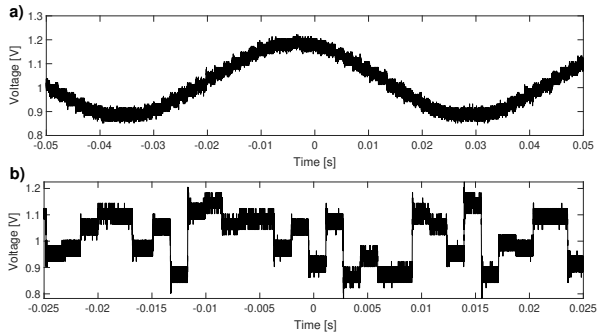

Fig. 4 Generated voltage waveforms using a) sine, and b) random function.

calculate the real heart rate and compare it with the heart rate obtained in the attack to test whether the attack is successful. The heart rate of the healthy heartbeats varies between 3 Hz and 3.2 Hz and for heartbeats with arrhythmia, between 3.6 Hz and 3.8 Hz.

The captured voltage over the resistor during one second for a Mat recording with healthy heartbeats is shown in Fig. 5a), while the TX signal is shown in Fig. 5b). The captured voltage has a lot of noise and the peaks due to the chip activity cannot be appreciated. The recorded signal needs to be processed to obtain the heart rate. After calculating the square of the voltage and applying median filter, the resulting signal finally reveals the heart rate as seen in Fig 5c) where the peak of the activity occurs slightly before the transmission of the fits seen in Fig. 5b). The median filter is applied with 10000 samples as it was observed that this number provides the most optimal results by manually adjusting the window value. For this window to work, it should include enough samples to cover the activity burst with a margin. For the oscilloscope set at 500 ms/division, the sampling rate is 400 kS/s, giving a window of 25 ms, enough to cover the activity burst but smaller than the 300 ms between peaks. Filtering the signal by using a Butterworth filter was discarded as it resulted in more noise in the signal, particularly in time domain. FFTs of the processed signal and TX line are calculated and shown in Fig. 5d), and the peaks at 3 Hz obviously match.

To test the attack success rate, we make 10 recordings for each configuration. We record the signal for 1 s, 2 s, and 5 s and we apply both scenarios to healthy heartbeats and heartbeats with arrhythmia. The results for the Mat attack are shown in Fig. 6 for recordings 100 and 201. FFT for the healthy heartbeats is marked in light grey, and for heartbeats with arrhythmia in black. FFTs for 1s recordings (Fig. 6a) seem to differentiate well between healthy and arrhythmia heartbeats as 3 Hz is detected for healthy and 4 Hz for heartbeats with arrhythmia. However, this is due to the poor frequency resolution as each frequency bin is 1 Hz wide and the true heart rates are rounded to the nearest integer. For 2 s recordings and 0.5 Hz frequency bin (Fig. 6b), the rate for heartbeats with arrhythmia oscillates between 3.5 Hz and 4 Hz for the same reason. The spectrum for 5 s recordings (Fig. 6c) presents anomalies due to the low sampling frequency that results in large peaks around 1.2 Hz for the healthy heartbeats and several recordings also have the heart rate equal to the heart rate of heartbeats with arrhythmia.

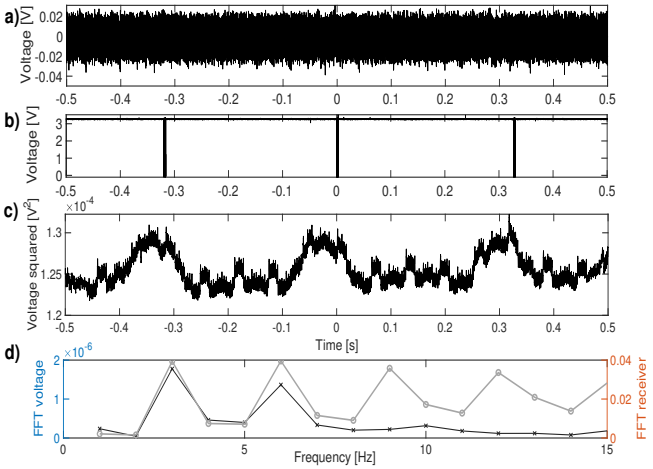

Fig. 5 a) Recorded voltage, b) transmission line signal, c) processed voltage in the time domain, d) FFT of the transmission line signal (light grey) and processed voltage (black)

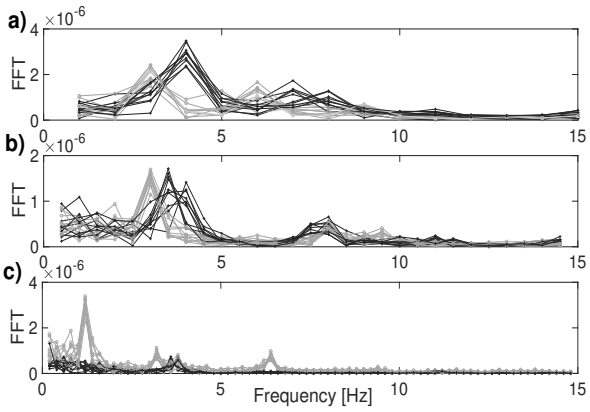

Fig. 6 FFT of the processed voltage for a) 1 s, b) 2 s, and c) 5 s recordings for healthy heartbeats (light grey) and heartbeats with arrhythmia (black).

Normally, when the ECG is recorded in real time, the resulting frequencies are in the [1 Hz, 1.7 Hz] range, and the difference between the normal and accelerated heart rates would be impossible to detect with any of the recording times tested here. The measurements with the probe that are recorded and saved on a USB key can only serve as an approximation of the heart rate, but not to differentiate between different heart rates.

Next, we apply Osc scenario to the healthy heartbeats and capture the resulting FFT as a bitmap image. The FFT

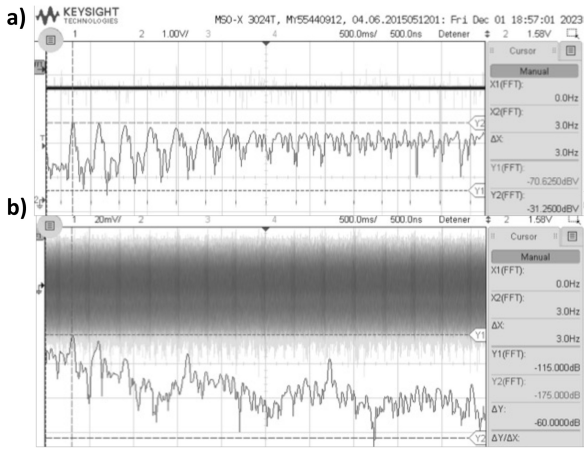

Fig. 7 Time domain and FFT image for (a) transmission signal. (b) processed voltage for healthy heartbeats.

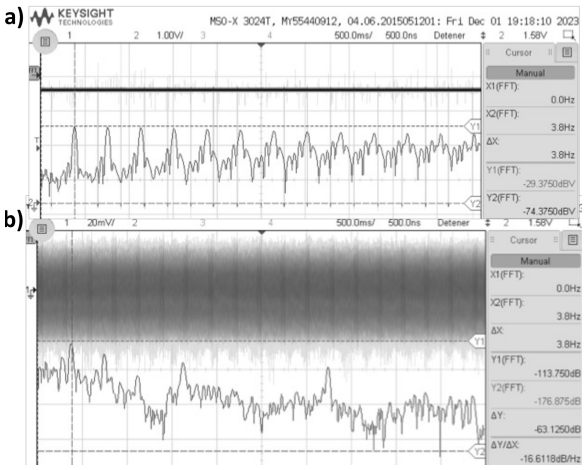

Fig. 8 Time domain and FFT image for (a) transmission signal. (b) processed voltage for heartbeats with arrhythmia.

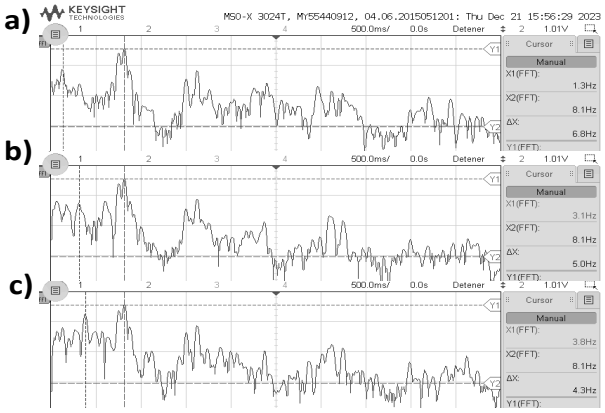

Fig. 9 FFT image for 8.1Hz sine modulation with (a) no data, (b) healthy heartbeats and c) heartbeats with arrhythmia.

for the TX line of recording 100 is shown in Fig. 7a) and for the measured voltage in Fig 7b). In each figure, the signal is presented both in time (above) and frequency (below). The X2 cursor is set to the heartbeat frequency. It can be seen that the 3 Hz peaks from Fig. 7a) and Fig. 7b) match. This peak remains visible in the voltage spectrum during the whole observation time. We also tested recordings 100, 101, 116, 119 and 220 and were able to obtain the heart rate by observing the frequency peak, which was set on frequencies between 2.6 Hz and 3 Hz.

To test whether the frequency of the accelerated heartbeats is captured correctly, we apply Osc scenario to the heartbeats with arrhythmia. The results are shown in Fig. 8 for recording 201. Due to the high sampling rate of the oscilloscope, the change in frequency between healthy and accelerated heartbeats is appreciated in the image. The dominant frequency is 3.8 Hz and is equal to the true frequency of arrhythmia heartbeats. We also tested recordings 108, 113, 115, and 123 and obtained their corresponding heart rate, thereby confirming the effectiveness of the proposed attack. The Osc attacks have higher success rate than Mat attacks, as the sampling frequency of the recordings used for the Matlab processing offers too little resolution for the activity peaks to be distinguished. Therefore, more accurate Osc attacks will be used to validate the proposed countermeasures.

The attacks could be further automatized by using high-speed ADC for power trace acquisition and Matlab for filtering and peak detection. Although they are carried out in lab conditions, the methodology could also be applied to electromagnetic leaks, not needing physical access.

4.2 Countermeasure results

Two different voltage modulations are applied: sine and random. The voltage range for both is [0.9 V, 1.15 V] to ensure correct functionality of the FPGA. The maximum frequency for the sine voltage is 15.6 Hz, a limitation due to the regulator time resolution. Several sine frequencies between 1 Hz and 15.6 Hz are applied during Osc attack. We record the leaked voltage spectrum when the application is idle and waiting for the data, when it is processing healthy heartbeats, and processing heartbeats with arrhythmia. For almost all frequencies, the heartbeat frequency peak in the spectrum has a smaller amplitude but is still clearly visible most of the time (see Fig. 9b) for healthy and Fig. 9c) for heartbeats with arrhythmia when the 8.1 Hz modulation frequency is applied). For two frequencies, 15.6 Hz and 13 Hz, the frequency peak is masked by the modulation as all three spectra are identical (see Fig. 10 for 15.6 Hz modulation frequency).

Even though applying sine modulation at these frequencies achieves heart rate masking, the peak at the modulation frequency is also clearly visible in the spectrum, and could also be deduced from the time domain signal, because a sine produces regular, repeatable oscillations. This information

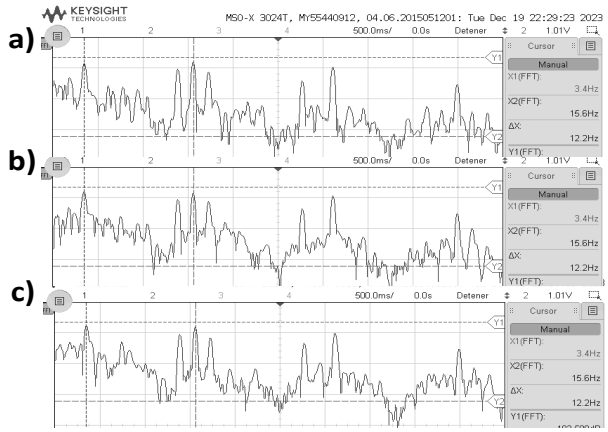

Fig. 10 FFT image for 15.6 Hz sine modulation with (a) no data, (b) healthy heartbeats and c) heartbeats with arrhythmia.

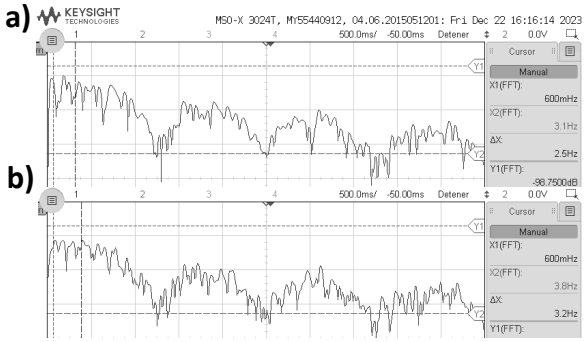

Fig. 11 FFT image for random modulation with 17.8 ms resolution for (a) healthy heartbeats and b) heartbeats with arrhythmia.

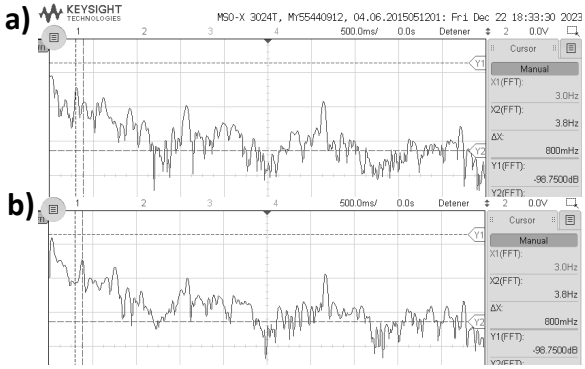

Fig. 12 FFT image for random modulation with 2.05s resolution for (a) healthy heartbeats and b) heartbeats with arrhythmia.

can be used by the attacker to demodulate the voltage estimating the frequency from the time series or spectrum and multiplying by another sine at the same frequency to obtain the victim's heart rate. For the random modulation, we tried different sample time resolutions from 1.6 ms to 10 s. The heart rate is completely masked for time resolutions below a 1.5s resolution, while it can be retrieved when the changes in the random voltage are above them. The results are presented in Fig. 11 and Fig. 12 for time resolutions of 17.8 ms and 2.05 s respectively. The random modulation is the most effective in hiding the processing inside the chip, but should be used with care since random spikes in the power supply can accelerate degradation of the device.

5 Conclusion

We present a methodology for power side-channel attacks on an ECG characterization prototype device. Two different equipment configurations are used to simulate an attacker's capabilities for obtaining heart rates from victims. The results show that the heart rate can be retrieved by processing directly the leaked power signal with the functions available in the oscilloscope. The signal that is recorded and then processed in Matlab has low resolution, due to the low number of saved samples and cannot be used to pinpoint the exact heart rate nor to distinguish the difference between the frequencies of healthy heartbeats and heartbeats with arrhythmia.

We then test the effectiveness of a countermeasure based on voltage modulation. Sine modulation at certain frequencies masks the heart rate, but the modulation frequency is also visible in the spectrum, enabling the attacker to demodulate the signal and retrieve the heart rate. Random modulation is effective against the attacks as it masks the heart rate completely when the resolution for the voltage change is less than 1.5 s. It should be used with caution, as abrupt voltage changes can reduce the durability of the device.

Acknowledgements The authors would like to thank Rodrigo Fernandez for his kind help recording power measurements.

Statements and Declarations

Funding No funding was received for conducting this study.

Competing interests The authors have no relevant financial or non-financial interests to disclose.

References

1. Cost of a Data Breach Report 2023. IBM Corporation (2023)
2. Sehatbakhsh, N., Yilmaz, B., Zajic, A., Prvulovic, M.: A new side-channel vulnerability on modern computers by exploiting electromagnetic emanations from the power management unit. In: In Proc. on IEEE Int. Symp. On High-Performance Computer Architecture (HPCA) (2020). DOI https://doi.org/10.1109/HPCA47549.2020.00020
3. Camurati, G., Poeplau, S., Muench, M., Hayes, T., Francillon, A.: Screaming channels: When electromagnetic side channels meet radio transceivers. In: In Proc. 2018 ACM SIGSAC Conference on Computer and Communications Security (2018). DOI https://doi.org/10.1145/3243734.3243802

4. Mangard, S., Oswald, E., Popp, T.: Power analysis attacks: Revealing the secrets of smart cards, vol. 31. Springer Science & Business Media (2008). DOI https://doi.org/10.1007/978-0-387-38162-6
5. Xu, R., Zhu, L., Wang, A., Du, X., Choo, K., Zhang, G., Gai, K.: Side-channel attack on a protected rfid card. IEEE Access **6** (2018). DOI https://doi.org/10.1109/ACCESS.2018.2870663
6. Genkin, D., Pipman, I., Tromer, E.: Get your hands off my laptop: physical side-channel key-extraction attacks on pcs: Extended version. J. Cryptogr. Eng. **5** (2015). DOI https://doi.org/10.1007/s13389-015-0100-7
7. Oswald, E., Mangard, S., Pramstaller, N., Rijmen, V.: A side-channel analysis resistant description of the AES S-Box. In: H. Gilbert, H. Handschuh (eds.) Fast Software Encryption. FSE 2005. Lecture Notes in Computer Science. Springer International Publishing (2005). DOI https://doi.org/10.1007/11502760_28
8. Banerjee, U., Ho, L., Koppula, S.: Power-based side-channel attack for AES key extraction on the ATMega328 microcontroller. arXiv (2022). DOI https://doi.org/10.48550/arXiv.2203.08220
9. Jevtic, R., Perez-Tirador, P., Cabezaolias, C., Carnero, P., Caffarena, G.: Side-channel attack countermeasure based on power supply modulation. In: 2022 30th European Signal Processing Conference (EUSIPCO), pp. 618–622 (2022). DOI https://doi.org/10.23919/EUSIPCO55093.2022.9909766
10. Singh, A., Kar, M., Mathew, S., Rajan, A., De, V., Mukhopadhyay, S.: Improved power/EM side-channel attack resistance of 128-bit AES engines with random fast voltage dithering. IEEE J. Solid-State Circuits **54**(2), 569–583 (2019). DOI https://doi.org/10.1109/JSSC.2018.2875112
11. Utyamishev, D., Partin-Vaisband, I.: Real-time detection of power analysis attacks by machine learning of power supply variations on-chip. IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. (2018). DOI https://doi.org/10.1109/TCAD.2018.2883971
12. Gattu, N., Imtiaz Khan, M.N., De, A., Ghosh, S.: Power side channel attack analysis and detection. In: IEEE/ACM International Conference On Computer Aided Design (ICCAD), pp. 1–7 (2020). DOI https://doi.org/10.1145/3400302.3415692
13. Pycroft, L., Aziz, T.: Security of implantable medical devices with wireless connections: The dangers of cyber-attacks. Expert Rev. Med. Devices **15**(6), 403–406 (2018). DOI https://doi.org/10.1080/17434440.2018.1483235
14. Liu, J., Homayoun, H., Fang, C., Miao, N., Wang, H.: Side channel-assisted inference attacks on machine learning-based ECG classification. In: 2023 IEEE/ACM Int. Conf. Computer Aided Design (ICCAD), pp. 1–9 (2023). DOI https://doi.org/10.1109/ICCAD57390.2023.10323617
15. ICS medical advisory (icsma-19-080-01), "Medtronic Conexus Radio Frequency Telemetry Protocol" (update C), cybersecurity and infrastructure security agency (2021)
16. Kune, D., Backes, J., Clark, S., Kramer, D., Reynolds, M., Fu, K., Kim, Y., Xu, W.: Ghost talk: Mitigating EMI signal injection attacks against analog sensors. In: 2013 IEEE Symposium on Security and Privacy, pp. 145–159 (2013). DOI https://doi.org/10.1109/SP.2013.20
17. Giechaskiel, I., Rasmussen, K.: Taxonomy and challenges of out-of-band signal injection attacks and defenses. IEEE Communications Surveys & Tutorials **22**(1), 645–670 (2020). DOI https://doi.org/10.1109/COMST.2019.2952858
18. Kim, Y., Lee, W., Raghunathan, V., Jha, N., Raghunathan, A.: Vibration-based secure side channel for medical devices. In: Proc. 52nd Annual Design Automation Conference, DAC '15. Association for Computing Machinery, New York, NY, USA (2015). DOI https://doi.org/10.1145/2744769.2744928
19. Di Cesare, M., et al.: World Heart Report 2023: Confronting the World's Number One Killer. World Heart Federation (2023)
20. Desai, M., Caffarena, G., Jevtic, R., Márquez, D., Otero, A.: A low-latency, low-power FPGA implementation of ECG signal characterization using Hermite polynomials. Electronics **10**(19) (2021). DOI https://doi.org/10.3390/electronics10192324
21. Lakhotia, K., Caffarena, G., Gil, A., Márquez, D., Otero, A., Desai, M.: Low-power, low-latency Hermite polynomial characterization of heartbeats using a field-programmable gate array. In: F. Ortuño, I. Rojas (eds.) Bioinformatics and Biomedical Eng., pp. 266–276. Springer International Publishing (2016). DOI https://doi.org/10.1007/978-3-319-31744-1_24
22. Chen, Z., Luo, J., Lin, K., Wu, J., Zhu, T., Xiang, X., Meng, J.: An energy-efficient ECG processor with weak-strong hybrid classifier for arrhythmia detection. IEEE Trans. Circuits and Systems II: Express Briefs **65**(7), 948–952 (2018). DOI https://doi.org/10.1109/TCSII.2017.2747596
23. Bock, C., Kovacs, P., Laguna, P., Meier, J., Huemer, M.: Ecg beat representation and delineation by means of variable projection. IEEE Trans. Biomed. Eng. **68**(10) (2021). DOI https://doi.org/10.1109/TBME.2021.3058781
24. Weimann, K., Conrad, T.: Transfer learning for ECG classification. Scientific Report 11 (2020). DOI https://doi.org/10.1038/s41598-021-84374-8
25. Aziz, S., Ahmed, S., Alouini, M.: ECG-based machine-learning algorithms for heartbeat classification. Scientific Report 11 (2021). DOI https://doi.org/10.1038/s41598-021-97118-5
26. Kovács, P., Fridli, S., Schipp, F.: Generalized rational variable projection with application in ecg compression. IEEE Trans. Signal Processing **68** (2020). DOI https://doi.org/1109/TSP.2019.2961234
27. Lagerholm, M., Peterson, C., Braccini, G., Edenbrandt, L., Sornmo, L.: Clustering ECG complexes using Hermite functions and self-organizing maps. IEEE Trans. Biomed. Eng. **47**(7) (2000). DOI https://doi.org/10.1109/10.846677
28. Haraldsson, H., Edenbrandt, L., M., O.: Detecting acute myocardial infarction in the 12-lead ecg using hermite expansions and neural networks. Artif. Intell. Med. **32**(2) (2004). DOI https://doi.org/10.1016/j.artmed.2004.01.003
29. Engelse, W., Zeelenberg, C.: A single scan algorithm for QRS-detection and feature extraction. Computers in cardiology **6**(1979), 37–42 (1979)
30. Adnane, M., Jiang, Z., Choi, S.: Development of QRS detection algorithm designed for wearable cardiorespiratory system. Computer Methods and Programs in Biomedicine **93**(1), 20–31 (2009). DOI https://doi.org/10.1016/j.cmpb.2008.07.010
31. Zidelmal, Z., Amirou, A., Adnane, M., Belouchrani, A.: QRS detection based on wavelet coefficients. Computer Methods and Programs in Biomedicine **107**(3), 490–496 (2012). DOI https://doi.org/10.1016/j.cmpb.2011.12.004
32. Li, Y., Yan, H., Hong, F., Song, J.: A new approach of QRS complex detection based on matched filtering and triangle character analysis. Australasian physical & engineering sciences in medicine **35**, 341–356 (2012). DOI https://doi.org/10.1007/s13246-012-0149-x
33. de Souza Faria, G., Kim, H.: Differential audio analysis: a new side-channel attack on PIN pads. Int. J. Inf. Secur. **18**, 73–84 (2019). DOI https://doi.org/10.1007/s10207-018-0403-7
34. Sonmez, B., Sarikaya, A., Bahtiyar, S.: Machine learning based side channel selection for time-driven cache attacks on AES. In: 2019 4th International Conference on Computer Science and Engineering (UBMK), pp. 1–5 (2019). DOI https://doi.org/10.1109/UBMK.2019.8907211
35. Camurati, G., Poeplau, S., Muench, M., Hayes, T., Francillon, A.: Screaming channels. In: Proc. of the 2018 ACM SIGSAC Conference on Computer and Communications Security. ACM (2018). DOI https://doi.org/10.1145/3243734.3243802
36. Spence, A., Bangay, S.: Security beyond cybersecurity: side-channel attacks against non-cyber systems and their countermeasures. Int. J. of Inf. Secur. **21**(3), 437–453 (2022). DOI https://doi.org/10.1007/s10207-021-00563-6

37. Spence, A., Bangay, S.: Side-channel sensing: Exploiting side-channels to extract information for medical diagnostics and monitoring. IEEE Journal of Translational Engineering in Health and Medicine **8**, 1–13 (2020). DOI https://doi.org/10.1109/JTEHM.2020.3028996
38. Jevtic, R., Otero, M.: Methodology for complete decorrelation of power supply EM side-channel signal and sensitive data. IEEE Trans. Circuits and Systems II: Express Briefs **69**(4), 2256–2260 (2022). DOI https://doi.org/10.1109/TCSII.2022.3144071
39. Jevtic, R., Ylitolva, M., Calonge, C., Ojanen, M., Santti, T., Koskinen, L.: EM side-channel countermeasure for switched-capacitor DC–DC converters based on amplitude modulation. IEEE Trans. Very Large Scale Integration (VLSI) Systems **29**(6), 1061–1072 (2021). DOI https://doi.org/10.1109/TVLSI.2021.3070687
40. Yu, W., Köse, S.: Charge-withheld converter-reshuffling: A countermeasure against power analysis attacks. IEEE Trans. Circuits and Systems II: Express Briefs **63**(5), 438–442 (2016). DOI https://doi.org/10.1109/TCSII.2015.2505261
41. Kamoun, N., Bossuet, L., Ghazel, A.: Correlated power noise generator as a low cost DPA countermeasures to secure hardware AES cipher. In: 3rd International Conference on Signals, Circuits and Systems (SCS), pp. 1–6 (2009). DOI https://doi.org/10.1109/ICSCS.2009.5412604
42. Parrilla, L., Garcia, A., Castillo, E., Rodriguez-Bolivar, S., Lopez-Villanueva, J.: Time- and amplitude-controlled power noise generator against SPA attacks for FPGA-based IoT devices. J. Low Power Electron. and Appl. **12**(3) (2022). DOI https://doi.org/10.3390/jlpea12030048
43. Newae Technology Inc.: CW305 Artix FPGA Target. URL https://www.newae.com/products/nae-cw305. Accessed 22nd February, 2024
44. Rinta-Aho, T., Karlstedt, M., Desai, M.: The Click2NetFPGA toolchain. In: 2012 USENIX Annual Technical Conference (USENIX ATC 12), pp. 77–88. USENIX Association, Boston, MA (2012)
45. Pololu USB-to-serial adapter. URL https://www.pololu.com/product/391. Accessed 22nd February, 2024
46. Moody, G., Mark, R.: The impact of the MIT-BIH arrhythmia database. IEEE Eng. Med. and Biol. Magazine **20**(3), 45–50 (2001). DOI https://doi.org/10.1109/51.932724
47. Texas Instruments TPS56520. URL https://www.ti.com/product/TPS56520. Accessed 9th July, 2024
48. Microchip Technology ATSAM3U2E. URL https://www.microchip.com/en-us/product/ATSAM3U2E. Accessed 9th July, 2024