



ThermWare: Toward Side-channel Defense for Tiny IoT Devices

Nakul Garg, Irtaza Shahid, Erin Avllazagaj, Jennie Hill, Jun Han[†], Nirupam Roy
University of Maryland College Park, [†]Yonsei University
{nakul22,irtaza,albocode,jehill}@umd.edu,jun.han@yonsei.ac.kr,niruroy@umd.edu

ABSTRACT

As malware in IoT devices flourishes, defenses are lacking. Traditional antivirus or intrusion detection-based defense techniques fail for the limited computational capabilities and the large diversity of platforms and environments. In this paper, we present *ThermWare*, a non-intrusive screening method to detect anomalous operations on embedded devices at run-time. *ThermWare* relies on the observation that electronic circuits generate subtle patterns of heat at the component level when the corresponding module is accessed by the micro-operations (e.g., file-write) of the running code. We propose the use of these side-channel heat signatures captured by a thermal camera to determine the sequence of underlying computations in real time. An early implementation of *ThermWare* shows success in detecting common malware routines in general-purpose IoT devices. We envision leveraging the thermal side-channel to track the internal operations of an embedded device, which can potentially lead to broader applications in engineering embedded systems, monitoring device health and run-time capacity, assisting embedded coding optimization, and physical layer security analysis.

CCS CONCEPTS

• Security and privacy → Security in hardware.

KEYWORDS

Malware detection; thermal side-channel; security; embedded systems

ACM Reference Format:

Nakul Garg, Irtaza Shahid, Erin Avllazagaj, Jennie Hill, Jun Han[†], Nirupam Roy. 2023. *ThermWare: Toward Side-channel Defense for Tiny IoT Devices*. In *The 24th International Workshop on Mobile Computing Systems and Applications (HotMobile '23)*, February 22–23, 2023, Newport Beach, CA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3572864.3580339>

1 INTRODUCTION

Miniaturized Internet of Things (IoT) devices have enabled a paradigm of connected everyday objects, ranging from a smart light bulb to implantable microchips for pets. Recent advancements show ubiquitous micro-motes [10] or insect-scale smart devices [7, 8, 19, 35, 47] for fine-grained sensing of our environments are not a distant future. The small size and the embedded nature of the future IoT devices forces a reconsideration of how we

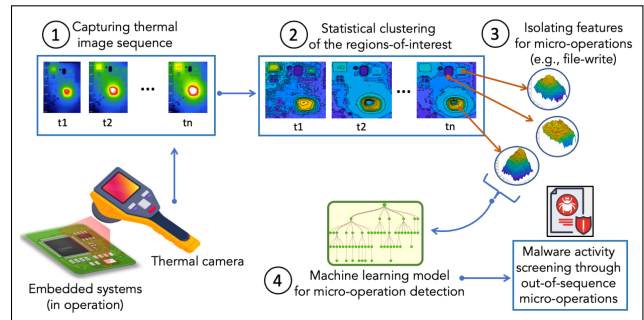


Figure 1: The overview of *ThermWare*, a thermal-side channel based malware detection system.

approach security – even the most basic defenses, like anti-virus software or process memory isolation, simply cannot run on resource-constrained systems. Often, hardware debugging ports are also sacrificed in the architecture in the interest of saving space and computation power on the circuit boards. As a result, standard security methods do not apply to these devices, which makes them notoriously vulnerable despite their close access to users’ private information. At the same time, it is difficult to make complete architectural changes to reinforce security in such systems [46]. One of the approaches to ensure security in resource-constrained embedded devices seeks alternative solutions in leveraging *side-channel* information for offloading anomaly detection, memory forensics, memory rewriting, and malware activity monitoring to more powerful machines.

The research community is constantly in search of new information sources to combine and complement side-channel defense techniques [20, 36, 42, 62]. This paper envisions a noninvasive malware detection approach using thermal maps. Activities on a processor and other peripherals lead to temperature variations at different regions of the circuit. The generated heat is directly related to the switching of transistors’ state at the lowest level – an indicator of activities at various functional units of the circuit and inside the chips. When monitored with fine-grained resolution in time and space, the transient thermal map of the circuit can reveal underlying operations at run-time. Such monitoring can be done passively with thermal cameras, without any physical contact with the device. A system can potentially learn this thermal pattern to detect any deviation from the desired activities on the circuit. We envision characterizing malicious activities on IoT devices through spatiotemporal analysis of the circuit’s heat dissipation. We call this system *ThermWare*. Figure 1 shows an overview of the system. If successful, this can become a noninvasive method for malware detection in small devices that are otherwise difficult to connect to standard monitoring equipment.

As a malware screening system, *ThermWare* requires only passive access to the circuit board for scanning thermal signatures using



This work is licensed under a Creative Commons Attribution International 4.0 License.

HotMobile '23, February 22–23, 2023, Newport Beach, CA, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0017-0/23/02.

<https://doi.org/10.1145/3572864.3580339>

infrared sensors or thermal cameras. Thermal map observation does not require physical or electrical contact with the device, which adds to its advantage compared to other side-channel methods. Our system can be used as an add-on to the device under monitoring and run detection code based on real-time thermal features. *ThermWare* can also be useful for auditing the behavior of small embedded devices, such as in-body medical devices when a patient visits the doctor’s office or during regular maintenance or replacements. While we believe the use-case and physical implementation of the method will evolve, in this paper we focus on exploring the core capabilities of this side-channel in malware detection and assess the advantages and limitations of such a system.

A system-on-chip (SoC) is built with careful positioning of modules like register banks, memory units, CPU, communication radio, and I/O interface inside one chip. We ask the question - *Can we leverage high-resolution thermal imaging techniques to detect anomalous computations inside a tiny SoC?* A program or computing operation can be viewed as a sequence of activation of these modules at different locations of the chip. Given such activation inevitably release heat, the spatial heat pattern can be a direct indicator of the underlying computing operations. This can open an opportunity to infer the executed code without a physical access. We aim to explore the possibility.

Detection of a micro-operation in run-time, such as file writes, is particularly relevant since the majority of malware in Linux environments are ransomware and web shells, and therefore file writes comprise a dominant portion of their behavior [6]. In this work, we show that we can detect file writes of at least 400KB with 80% accuracy using a relatively inexpensive thermal camera. Additionally, we achieve more than 90% accuracy with a higher-end thermal camera. Our results suggest that a high-end camera provides diminishing returns if an application tolerates a slight decrease in the accuracy level. Next, we show that an anomaly detector can achieve reliable detection of common malware routines in general-purpose IoT devices using an inexpensive thermal camera adapter to a smartphone. At the same time, we report a few undetectable actions, such as network port scanning on a voice assistant device.

ThermWare demonstrates the feasibility of using heat signatures as an alternative method to detect the states of computational operations inside an embedded device. We hypothesize that the fundamental units of the computations, such as file write events, memory access, network operations, etc., manifest as thermal signatures. If this side-channel information is carefully analyzed, it can track the behavior of running codes. This technique can have several implications by adding to the capabilities of reverse engineering embedded systems, monitoring device health and run-time capacity, assisting embedded coding optimization, and physical layer security analysis. Needless to say, this paper is only a first step toward this broader vision and we aim to inquire about the opportunities it presents. We start by experimenting with the technique as a defense application that can identify anomalous behavior and malware in an embedded system without requiring system-level access.

2 FEASIBILITY STUDY

For initial proof-of-testing, we focus on detecting malware running on a single-board-computer. While many operations could be used as features to identify malware, we chose to specifically focus on file writes for this initial work, as malware most commonly performs file writes [6]. By extension, if we can identify the presence of erroneous file writes, we can detect the presence of malware.

We perform a 100MB file write on a BeagleBone Black [1] and record the thermal images using a Seek thermal camera [2] at 9 fps. Figure 2 shows the board with three components indicated, as well as a series of thermal images of the board’s heat signature over the course of a file write. All three chips provide clear visual indications of write activity, as each gets noticeably hotter over milliseconds. The eMMC flash storage is of particular interest as it only heats up during file writes, lending credibility to the hypothesis that the thermal signature could be used to identify the presence of malware.

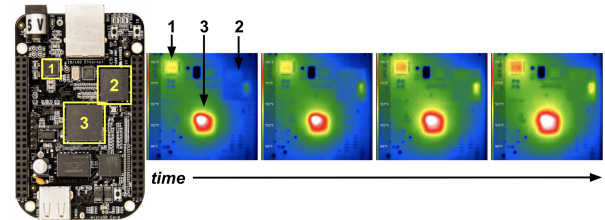


Figure 2: Thermal signature of a 100 MB file written over time. The left image shows a BeagleBone Black with three components highlighted: (1) Power Management Integrated Chip, (2) Embedded Multimedia Card, and (3) Processor, all of which produce distinct heat signatures over the course of the write operation.

Complimenting power side-channel: Prior work [14] has shown the feasibility of using a power side channel to detect generic malware in the wild. We analyze the difference between power and thermal side channels. In Figure 3 we show the power consumption and average pixel value of the power management chip of Beaglebone while running a Gafgyt malware. Formally, we measured that the 2 signals are strongly correlated with a 0.76 Pearson correlation coefficient. However, unlike a power side channel, the thermal side channel is delayed by around 0.93s. This is because the chip takes time to heat, even though it immediately starts drawing power. Therefore, we consider the thermal side channel as a complementary set of features to the power side channel, for the task of malware detection.

3 DETECTING FILE-WRITE OPERATIONS FOR MALWARE DETECTION

File write operations involve writing some number of bytes into the onboard solid-state drive (SSD)/eMMC of the embedded system. This writing activity causes the SSD chip to heat up, as shown previously in Figure 2. To develop a malware detector, we focus on detecting these heat variations in the solid-state drive. In this section, we will describe a manual image processing-based and learning-based technique to localize temperature variations and classify if a file write operation is performed.

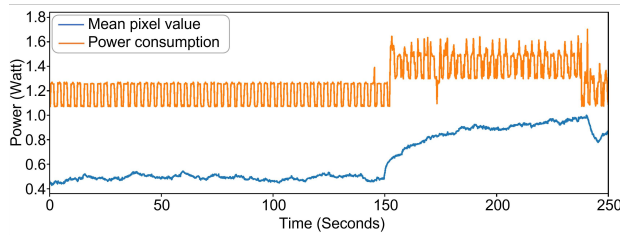


Figure 3: Power vs. thermal side channel when running Gafgyt botnet on a Beaglebone Black.

3.1 Image Processing-based Detection

We process the recorded video frames individually to extract the features needed for classification. As the first work in this area, we chose the most common features: average pixel intensity and the ratio of heated area. The former is calculated as the mean value of the pixels in the image and the second is the number of pixels above a certain value. This step yields two values for each frame of the video. As a next step, we merge the results of the calculation for each frame into a time series. To remove the noise, we use a Butterworth low-pass filter. We manually tuned the hyperparameters to remove most of the noise without impacting small file writes. We found that a filter order of 3 and a cutoff frequency of 0.1 yielded the best results for all the file writes. Figure 4 and 5 shows the mean pixel intensity values before and after filtering. We notice that the sharpest peaks correspond to actual file write operations, so we take the derivative of the time series signal. We put a threshold on the derivative to determine when a peak is due to an operation. The line in orange corresponds to part of the plot above the 0.001 threshold, which indicates a detected file write.

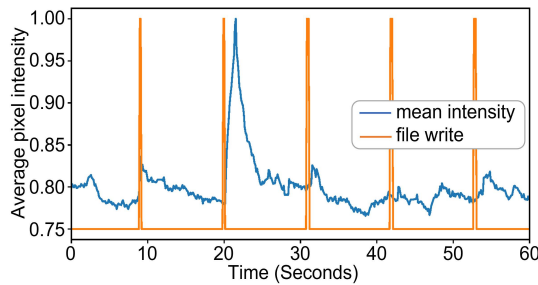


Figure 4: The extracted signal from the thermal camera and ground truth of file-write events over time.

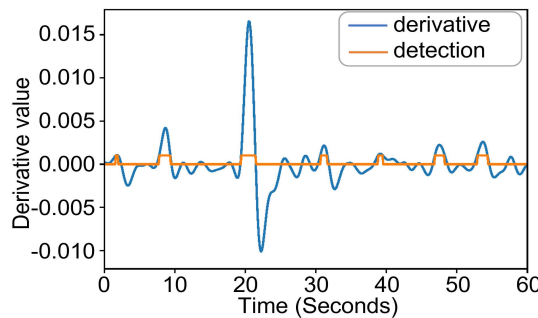


Figure 5: Filtered signal and detected anomalous file-writes events.

3.2 Learning-based Detection

We also explore a machine-learning-based approach to develop a general-purpose pipeline that can detect anomalies on any chip. We try Isolation Forests (IF) and Local Outlier Factors (LOF). After multiple trials, we concluded that for the correct hyperparameter values IF and LOF would produce the same anomalies. Therefore, we chose IF for our system, since it requires less memory for inference.

Feature extraction: For a general approach, independent of the type of Integrated Circuit, we first find the region of interest. We use a time series of all the values each pixel takes over time and calculate the min-max variability. The intuition is that the pixels with the largest variability are the pixels that heat up during an operation. We found that taking the top 3000 pixels gave us the best representation of the region in which we could visually observe a temperature increase. After slicing 3000 pixels from every video frame, we use the 9000 frames from the idle video to train an isolation forest with 50 trees. After training, we tested using the cropped video frames from the recording with file writes and got an anomaly score for each video frame. Finally, we use a confidence threshold to classify frames with anomalies.

Classification: We merge consecutive detected frames and those with a distance of less than 4 undetected frames (~ 0.5 seconds) into a *detection interval* and count how many of the intervals overlap with the ground truth of when a file write occurred. We remove any interval that has less than 5 frames because they are likely false positives. Since there is a delay from the occurrence of a file write to when the heat from the write is discernible from the SSD, we consider a ± 2 second interval from the operation was recorded as having been performed. Figure 6 shows the classification results along with the extracted signal from the region-of-interest. As we can see, the anomaly detector misses just 1 file write operation.

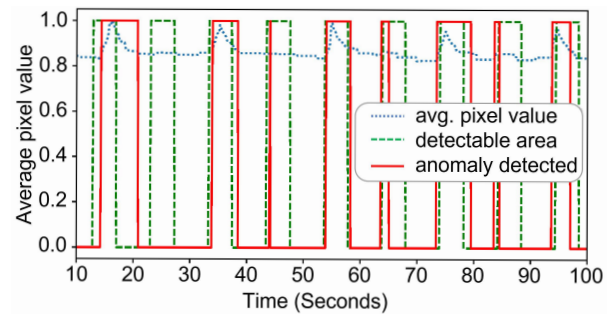


Figure 6: Detecting anomalous file write operations using the machine-learning-based technique.

4 IMPLEMENTATION AND SETUP

In this section, we describe our experimental setup and various IoT device attack scenarios.

4.1 Experimental Setup

The components involved in generating, collecting, and analyzing thermal side-channel data are described as follows:

1) **Single-Board Computer (SBC):** The experiments are performed on a BeagleBone Black SBC [1]. The SBC contains a

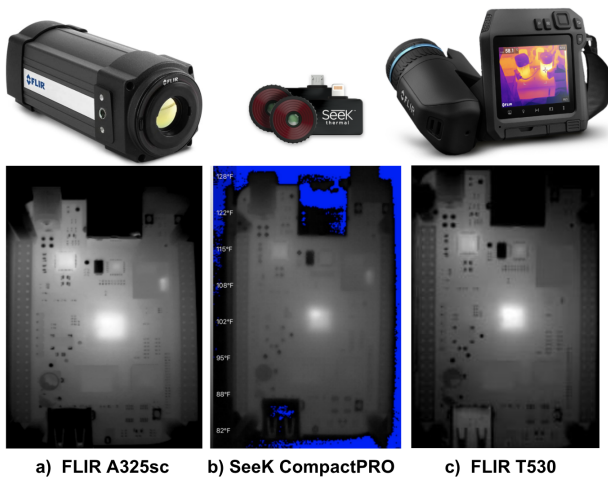


Figure 7: Comparison of images captured with (a) FLIR A325sc, (b) Seek Thermal CompactPRO, and (c) FLIR T530 cameras.

power manager, 4GB 8-bit embedded Multi-Media Card (eMMC), a Sitara AM3358 ARM processor, and 512 MB DDR3 RAM. It also runs the default Ubuntu 18.04 along with the GNU C and C++ compilers and Python 2.7.18. The SBC is controlled through a series of Python scripts to write files of sizes varying from 25KB to 1000KB to the BeagleBone’s flash storage, from a laptop.

2) Thermal Cameras: Two of the three thermal cameras shown in Figure 7: FLIR A325sc [3] and Seek CompactPRO [2] were chosen to collect data, due to their diversity in thermal sensitivity, frame rate, and price point. The detailed characteristics of these cameras are mentioned in Table 1. The FLIR A325sc has a thermal sensitivity of $50mK$ and is capable of capturing video at a rate of 60 frames per second. In contrast, the Seek CompactPRO has a slightly higher thermal sensitivity of $70mK$, and can only capture video at a rate of 9 frames per second. There is a significant price difference between these two cameras, with the FLIR A325sc costing \$10,000 and the Seek CompactPRO costing \$500. FLIR A325sc is connected to the laptop running FLIR Research Studio version 1.1.3 via ethernet cable. Seek CompactPRO is connected to the iPhone via a lightning cable and controlled through the official Seek Thermal application version 2.2.7.0. To maintain a stable video feed, the thermal camera is placed in a fixed position, so that the majority of the board is within the camera’s field of view. We ensure that the camera is not in contact with any component on the board.

| | FLIR A325sc | Seek CompactPRO | FLIR T530 |
|--------------------------|-------------|-----------------|------------|
| Resolution (pixels) | 320x240 | 320x240 | 320x240 |
| Temperature (°C) | -20 to 350 | -40 to 330 | -20 to 650 |
| Thermal Sensitivity (mK) | 50 | <70 | <50 |
| Frame Rate (Hz) | 60 | 9 | 30 |
| Price (USD) | 10K | 500 | 11K |

Table 1: Comparison of the thermal cameras used in experiments.

4.2 IoT Device Attack Scenarios

For testing different applications of IoT devices, we simulate a temperature sensor, air quality sensor, and a voice assistant on the Beaglebone Black.

•**Temperature and air quality sensors:** A simulation of the temperature and air quality sensor is performed by having the Beaglebone read `/dev/urandom` and log the values in a file, then send it through the network to a file server.

•**Voice assistants:** We simulate a voice assistant by recording audio of the top 6 Alexa commands and translating them to text using a speech-to-text library. We encompass a large number of commands that deal with data fetching from the internet and audio playback. We accomplish this with a script that downloads 10KB of data for 1000 iterations, performing text-to-speech and writing the results to a file.

•**Malware:** Malware is simulated while all the devices perform their usual tasks. We consider either idle or logging states for temperature and air quality sensors. For a voice assistant, we consider the following set: idle, getting voice commands, and playing music. We extract our features every 1 minute, such that we can capture all states. Our list of simulated malware actions is based on the top 4 most common Linux malware families [34]. We extract the SYN-flood routine from lightaidra, the telnet brute force from Mirai, network scanning from pnsnscan, and a generic CPU cryptomining program.

•**Data collection:** From the thermal recordings, every 1 minute we extract the following features: number of file writes detected by previous anomaly detector (1 feature), general statistics about CPU, power management chip, SSD and RAM (Mean, Variance, IQR, Skewness, Kurtosis, Min, Max) ($4 * 7 = 28$ features), and every 10 seconds we get 25th, 50th, 75th percentile of textural features [25] (contrast, dissimilarity, homogeneity) for distance 5 and angle 0 ($6 * 3 = 18$ features).

5 PRELIMINARY RESULTS

In this section, we evaluate the performance of our file write detectors and malware detection in IoT devices.

5.1 Performance: File Write Detection

We perform a systematic study in detecting file writes of different sizes. We perform 20 file writes for file sizes from 1000KB - 25KB. These files’ write event sizes are chosen to be equal to or less than the average malware file size (in the range of hundreds of KB) [40]; in fact, Lua malware are even larger than 1MB. We capture thermal video with both the FLIR A325sc and the Seek CompactPRO cameras. We then apply both our automatic and manual file write detection techniques from Sections 3.1 and 3.2, respectively.

Figure 8 shows the performance for file write detection using both types of cameras. We see that FLIR thermal camera has the highest performance at file sizes $> 300KB$, and is at or above 90% at file sizes of 500KB and above. Conversely, the Seek camera, averaged around 70% accuracy at file sizes $> 300KB$ and outperformed the FLIR at

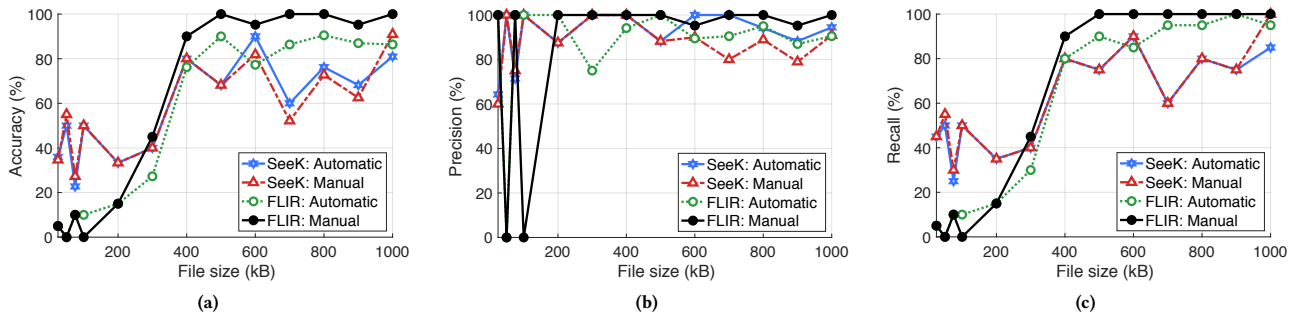


Figure 8: (a) Accuracy, (b) precision, and (c) recall of SeeK Compact Pro and FLIR A325sc cameras, using both automated and manual detection.

detecting file writes $\leq 300\text{KB}$, with accuracy ranging between roughly 30 and 50%. There’s a clear trade-off compared to the much less expensive SeeK if slightly lower detection performance is tolerable. Manual write detection slightly outperformed our automated write detector in both accuracy and recall, although additional evaluation of the automated procedure or alternate ML techniques would likely narrow the margin.

5.2 Performance: Malware Detection

In this section, we explore the capability of our features to detect malware in embedded devices. A detection is considered successful if the anomaly detector raises an alert within 1 minute of when the malware execution takes place. Similarly, a false positive is counted when an alert is raised during a time window where there is no malware running and a false negative is considered when no alert is raised during malware execution.

Table 2 shows the results of detecting various malicious routines as they run alongside the emulated devices. As mentioned, we run each of the malware routines 10 times for each benign action. Following are the key observations: 1) FP is usually 1 for Ransomware and Cryptomining activities. The reason is that these activities significantly heat the CPU, which takes some time to cool down. This causes our system to register anomalies even after malware activity is finished. 2) Our system finds it challenging to detect malware attacks that require very little CPU activity (e.g., scanning). For example, our system missed the first two runs of telnet brute force, when the scanner was searching for IPs to attack. 3) In the voice assistant case, the model is trained to accept long actions with high CPU requirements (e.g., playing music). Therefore, no alarm will be raised unless the CPU or power management chips heat beyond the detection threshold. That is why our detector missed some quick low computational network port scanning attacks. 4) It also highlights the limitation that a targeted malware may be able to evade our detector by operating in sequences, allowing chips to cool down before performing the next batch of operations.

6 RELATED WORK

The literature in embedded systems security and side-channel information processing is extremely rich and rooted in many subfields of engineering. Here we zoom into some related areas.

Securing IoT devices: The susceptibility of IoT devices to attack has been studied thoroughly [30, 51, 52, 54, 55, 58, 62]. In 2016, the

| | Device Action | Malware Action | TP | FP | FN |
|-----------------|----------------|----------------|----|----|----|
| Smart therm. | Idle | Ransomware | 10 | 1 | 0 |
| | | Cryptomining | 10 | 1 | 0 |
| | | Others | 20 | 0 | 0 |
| | Logging | Net Scanning | 9 | 0 | 1 |
| | | Others | 30 | 0 | 0 |
| Air sensor | Idle | Ransomware | 10 | 1 | 0 |
| | | Others | 30 | 0 | 0 |
| | Logging | Ransomware | 10 | 1 | 0 |
| | | Cryptomining | 10 | 1 | 0 |
| | | Others | 20 | 0 | 0 |
| Voice Assistant | Idle | All | 40 | 0 | 0 |
| | Speech-to-Text | Ransomware | 10 | 1 | 0 |
| | | Net Scanning | 9 | 0 | 1 |
| | | Others | 20 | 0 | 0 |
| | Play music | Ransomware | 10 | 1 | 0 |
| | | Telnet Brute. | 8 | 0 | 2 |
| Net Scanning | | 2 | 0 | 8 | |
| | | Cryptomining | 9 | 0 | 1 |

Table 2: Malware detection results for different kinds of devices and activities.

Mirai botnet was used to launch some of the largest recorded DDoS attacks in history, including a 623 Gbps attack [5], and a 1.2 Tbps attack [17]. A Mirai bot simply spread by attempting log into a Telnet service using a list of factory default username/password combinations. The type of solutions that have been considered thus far focuses predominately on analyzing the network traffic of these devices [9, 22, 38, 60, 62]. *ThermWare* is complementary to these methods and provides an independent source of detection and validation.

Attack using side-channels: Past projects have shown methods for inferring information from a variety of signal sources [21], including electromagnetic fields [4, 45], household AC power consumption [26], sounds and vibrations [12, 20, 49, 53], inertial sensor data [15], and optical channels [56, 59]. Integrated circuit (IC) design and manufacturing research have explored on-chip heat distributions as one of the quality metrics [24, 32, 39]. Security researchers have leveraged this as a side channel to probe internal activities and infer various information to reveal potential attacks through covert communications and data leakage [23, 29, 37, 41, 44]. Prior work [57] explores a reverse problem and has shown that CPU workload can be used to infer CPU temperature. In this paper,

we seek to utilize the inherent problem of information leakage to our advantage in detecting malware.

Defense using side-channels: The use of side channels to detect malicious behavior has also been prevalent in several prior works [16, 31, 43, 61]. In [11], side-channel current analysis is used to detect the presence of modified firmware on an embedded solid-state drive. A few early papers [13, 14, 33, 36], published between 2008 and 2013, have shown the possibility of using energy consumption patterns as an indicator of anomalous activity on computing devices. VirusMeter [36] and WattsUpDoc [14] show changes in power consumption during the execution of malware codes in cell phones and medical devices. Prior works [18, 27, 28] have used the thermal and EM side channels to detect backdoors in device manufacturing. While their threat model is different from ours (e.g., supply chain attacker), they demonstrate that these side channels can be used to detect anomalies in a chip’s normal heating patterns. Our work demonstrates that thermal readings can be used to detect anomalous out-of-order operations for the purpose of detecting malware.

7 APPLICATIONS BEYOND SECURITY

In this paper, we explored the potential of using thermal side-channel techniques to detect malware in embedded systems. We believe that *ThermWare* can also be extended to more general use cases and application scenarios such as non-intrusive reverse engineering and remote performance auditing.

Non-intrusive reverse engineering: Reverse engineering embedded systems can be a challenging and time-consuming process, especially if the goal is to do so in a non-intrusive way. However, the thermal side-channel can be a useful tool to gain insight in the internals of the chip. By analyzing the changes in surface temperature of the chip, as it performs different operations, an engineer can create a spatio-temporal heatmap which can reveal the sequence of operations that were executed. This information can be helpful in understanding the design and function of the system and can help in reverse engineering the ongoing process.

Performance monitoring: *ThermWare* can also be used to remotely monitor and audit the performance of data centers and high performance computing (HPC) clusters. The thermal signatures can be obtained using the existing temperature sensors and infrared cameras in the centers. These sensors can be repurposed to create a more fine grained spatial heatmap which can be then used to monitor the performance and behaviour of the system, including memory transfer, job scheduling and activity of CPU, GPU, routers, etc. Using the performance metrics, engineers can gain insights into the functioning of the data centers and identify potential issues or scope of improvements. Additionally, any non-regular temperature patterns detected can serve as early warnings for data leaks, malware, or other security threats.

Covert communication: Alternative communication modalities are useful [48, 50] for clandestine operations where information need to be exchanged without leaving any trace in the wireless channel or in the Internet. Recently, *BitWhisper* showed the use of heat generated by CPU or GPU as a modality for

covert data communication [23]. By detecting the heat signature of micro-operation at the circuit level, *ThermWare* essentially provides a way to observe complex patterns of heat signatures. If explored, this can be a building block for a high data-rate thermal communication channel between an embedded device and a thermal camera. The system can run a combination of micro-operations to generate a specific pattern of heat across different components encoding data bits. The camera can use methods shown in this paper to recover data from the temporal variations of the heat patterns.

8 CONCLUSION AND FUTURE WORK

This paper presents *ThermWare*, a malware detection system for IoT devices using thermal side-channel. We show the possibility of real-time malware detection by monitoring the thermal patterns of a file write operation on an embedded system. Needless to say, *ThermWare* is an exploratory first realization of the concept and there is room for improvements and further work.

Formalizing the verification process: *ThermWare* requires knowing the original thermal signature of an application to detect anomalies. This limits the generalization to other applications and software updates of the existing application since any change in code affects the thermal signature. We plan to learn the mapping of code to thermal signatures. Our intuition is to have a directed graph of computing operation to a spatio-temporal thermal sequence which can be used to generate new thermal signatures given any code. If successful, developers can create a version of true thermal signatures using the directed graph and release it along with their code.

Generalize malware detection: *ThermWare* is a novel solution for detecting malware that leverages the inherent temperature variations inside a chip that occur when executing commands. At present, the system is primarily focused on read/write activities. However, the fundamental concept of detecting malware by monitoring temperature patterns is not restricted to read/write activities alone. In the future, we aim to generalize this technique for the detection of any unauthorized execution of commands, as any process or computation will result in heat variations in certain parts of the chip. These variations may be small or large, depending on the activity, but with finely calibrated temperature sensors, we can detect these variations and use them to detect malware.

Obfuscation: Currently, the system does not account for an adversary that is invested in bypassing our thermal detection. For example, an adversary can obfuscate the attack by emulating thermal signatures. This can be done by carefully timing the file write operation along with other file write operations or by emulating a known application’s signature. We plan to work on identifying such obfuscation attacks.

Resolution: In the future, we will explore the effect of the spatial and temporal spread of the heat to neighboring regions. If a compute-heavy task takes place, the heat gradually spread all over the chip spoiling the resolution until the chip cools down. We want to make the system robust to past operations by incorporating the effect of the multidimensional Point Spread Function of the heat. As this research evolves, we envision that our core idea of thermal side-channel will open new ways of defense in miniature IoT devices.

REFERENCES

- [1] Beaglebone black. <https://beagleboard.org/black>. Last accessed 21 December 2022.
- [2] Compactpro: affordable, high-performance thermal imaging for your smartphone. <https://www.thermal.com/uploads/1/0/1/3/101388544/compactprosheet-usav1.pdf>. Last accessed 21 December 2022.
- [3] Flir a325sc thermal imaging camera for real-time analysis. https://www.flirmedia.com/MMC/THG/Brochures/RND_010/RND_010_US.pdf. Last accessed 21 December 2022.
- [4] AGRAWAL, D., ARCHAMBEAULT, B., RAO, J. R., AND ROHATGI, P. The em side-channel(s). In *Cryptographic Hardware and Embedded Systems - CHES 2002* (Berlin, Heidelberg, 2003), B. S. Kaliski, ç. K. Koç, and C. Paar, Eds., Springer Berlin Heidelberg, pp. 29–45.
- [5] AKAMAI. Akamai's State of the Internet / Security, Q3 2016 Report. <https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/q3-2016-state-of-the-internet-security-report.pdf>, 2016.
- [6] AVLLAZAGAJ, E., ZHU, Z., BILGE, L., BALZAROTTI, D., AND DUMITRAS, T. When malware changed its mind: An empirical study of variable program behaviors in the real world. In *30th USENIX Security Symposium (USENIX Security 21)* (Aug. 2021), USENIX Association, pp. 3487–3504.
- [7] BAI, Y., GARG, N., AND ROY, N. Spidr: Ultra-low-power acoustic spatial sensing for micro-robot navigation. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services (2022)*, pp. 99–113.
- [8] BAI, Y., GARG, N., AND ROY, N. Ultra-low-power acoustic imaging. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services (2022)*, pp. 523–524.
- [9] BENSON, T., AND CHANDRASEKARAN, B. Sounding the bell for improving Internet (of Things) security. In *ACM Workshop on Internet of Things Security and Privacy (IoTS&P)* (2017).
- [10] BICK, C. S., LEE, I., COOTE, T., HAPONSKI, A. E., BLAAUW, D., AND FOIGHIL, D. Ó. Millimeter-sized smart sensors reveal that a solar refuge protects tree snail partula hyalina from extirpation. *Communications Biology* 4, 1 (2021), 1–8.
- [11] BROWN, D., WALKER III, T. O., BLANCO, J. A., IVES, R. W., NGO, H. T., SHEY, J., AND RAKVIC, R. Detecting firmware modification on solid state drives via current draw analysis. *Computers & Security* (2020), 102149.
- [12] CHOUDHURY, R. R., AND ROY, N. Vibrational devices as sound sensors, Apr. 21 2020. US Patent 10,628,484.
- [13] CLARK, S. S., MUSTAFA, H., RANSFORD, B., SORBER, J., FU, K., AND XU, W. Current events: Identifying webpages by tapping the electrical outlet. In *European Symposium on Research in Computer Security* (2013), Springer, pp. 700–717.
- [14] CLARK, S. S., RANSFORD, B., RAHMATI, A., GUINEAU, S., SORBER, J., XU, W., AND FU, K. Wattspdoc: Power side channels to nonintrusively discover untargeted malware on embedded medical devices. In *Presented as part of the 2013 {USENIX} Workshop on Health Information Technologies* (2013).
- [15] DEY, S., ROY, N., XU, W., CHOUDHURY, R. R., AND NELAKUDITI, S. Accelprint: Imperfections of accelerometers make smartphones trackable. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)* (2014).
- [16] DING, F., LI, H., LUO, F., HU, H., CHENG, L., XIAO, H., AND GE, R. Deeppower: Non-intrusive and deep learning-based detection of iot malware using power side channels. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security* (New York, NY, USA, 2020), ASIA CCS '20, Association for Computing Machinery, p. 33–46.
- [17] DYN. Dyn analysis summary of Friday October 21 attack. <https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/>, 2016.
- [18] FORTE, D., BAO, C., AND SRIVASTAVA, A. Temperature tracking: An innovative run-time approach for hardware trojan detection. In *2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)* (2013), IEEE, pp. 532–539.
- [19] GARG, N., AND ROY, N. Enabling self-defense in small drones. In *Proceedings of the 21st International Workshop on Mobile Computing Systems and Applications* (2020), pp. 15–20.
- [20] GENKIN, D., SHAMIR, A., AND TROMER, E. Rsa key extraction via low-bandwidth acoustic cryptanalysis. In *Advances in Cryptology – CRYPTO 2014* (Berlin, Heidelberg, 2014), J. A. Garay and R. Gennaro, Eds., Springer Berlin Heidelberg, pp. 444–461.
- [21] GU, P., STOW, D., BARNES, R., KURSUN, E., AND XIE, Y. Thermal-aware 3d design for side-channel information leakage. In *2016 IEEE 34th International Conference on Computer Design (ICCD)* (2016), IEEE, pp. 520–527.
- [22] GUO, H., AND HEIDEMANN, J. IP-based IoT device detection. In *ACM Workshop on Internet of Things Security and Privacy (IoTS&P)* (2018).
- [23] GURI, M., MONITZ, M., MIRSKI, Y., AND ELOVICI, Y. Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations. In *2015 IEEE 28th Computer Security Foundations Symposium* (2015), IEEE, pp. 276–289.
- [24] HAMANN, H. F., LACEY, J., WEGER, A., AND WAKIL, J. Spatially-resolved imaging of microprocessor power (simp): hotspots in microprocessors. In *Thermal and Thermomechanical Proceedings 10th Intersociety Conference on Phenomena in Electronics Systems, 2006. IThERM 2006.* (2006), IEEE, pp. 5–pp.
- [25] HARALICK, R. M., SHANMUGAM, K., AND DINSTEN, I. Textural features for image classification. *IEEE Transactions on Systems, Man, and Cybernetics SMC-3*, 6 (1973), 610–621.
- [26] HART, G. W. Nonintrusive appliance load monitoring. *Proceedings of the IEEE* 80, 12 (1992), 1870–1891.
- [27] HE, J., GUO, X., MA, H., LIU, Y., ZHAO, Y., AND JIN, Y. Runtime trust evaluation and hardware trojan detection using on-chip em sensors. In *2020 57th ACM/IEEE Design Automation Conference (DAC)* (2020), IEEE, pp. 1–6.
- [28] HU, K., NOWROZ, A. N., REDA, S., AND KOUSHANFAR, F. High-sensitivity hardware trojan detection using multimodal characterization. In *2013 Design, Automation & Test in Europe Conference & Exhibition (DATE)* (2013), IEEE, pp. 1271–1276.
- [29] HUTTER, M., AND SCHMIDT, J.-M. The temperature side channel and heating fault attacks. In *International Conference on Smart Card Research and Advanced Applications* (2013), Springer, pp. 219–235.
- [30] JIA, Y. J., CHEN, Q. A., WANG, S., RAHMATI, A., FERNANDES, E., MAO, Z. M., AND PRAKASH, A. Contextiot: Towards providing contextual integrity to appified IoT platforms. In *Network and Distributed System Security Symposium (NDSS)* (2017).
- [31] JIMÉNEZ, J. M. H., NICHOLS, J. A., GOSEVA-POPSTOJANOVA, K., PROWELL, S., AND BRIDGES, R. A. Malware detection on general-purpose computers using power consumption monitoring: A proof of concept and case study. *arXiv preprint arXiv:1705.01977* (2017).
- [32] KENDIG, D., YAZAWA, K., MARCONNET, A., ASHEGHI, M., AND SHAKOURI, A. Side-by-side comparison between infrared and thermorefectance imaging using a thermal test chip with embedded diode temperature sensors. In *2012 28th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)* (2012), IEEE, pp. 344–347.
- [33] KIM, H., SMITH, J., AND SHIN, K. G. Detecting energy-greedy anomalies and mobile malware variants. In *Proceedings of the 6th international conference on mobile systems, applications, and services* (2008), pp. 239–252.
- [34] KINGER, P., AND LOGAN, M. Linux threat report 2021 1h: Linux threats in the cloud and security recommendations, Aug 2021.
- [35] LEE, I., HSIAO, R., CARICHNER, G., HSU, C.-W., YANG, M., SHOOURI, S., ERNST, K., CARICHNER, T., LI, Y., LIM, J., ET AL. msail: milligram-scale multi-modal sensor platform for monarch butterfly migration tracking. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking* (2021), pp. 517–530.
- [36] LIU, L., YAN, G., ZHANG, X., AND CHEN, S. Virusmeter: Preventing your cellphone from spies. In *International Workshop on Recent Advances in Intrusion Detection* (2009), Springer, pp. 244–264.
- [37] MASTI, R. J., RAI, D., RANGANATHAN, A., MÜLLER, C., THIELE, L., AND CAPKUN, S. Thermal covert channels on multi-core platforms. In *24th USENIX security symposium (USENIX security 15)* (2015), pp. 865–880.
- [38] MEIDAN, Y., BOHADANA, M., SHABTAI, A., OCHOA, M., TIPPENHAUER, N. O., GUARNIZO, J. D., AND ELOVICI, Y. Detection of unauthorized IoT devices using machine learning techniques. <https://arxiv.org/pdf/1709.04647.pdf>.
- [39] MESA-MARTINEZ, F. J., BROWN, M., NAYFACH-BATTILANA, J., AND RENUA, J. Measuring power and temperature from real processors. In *2008 IEEE International Symposium on Parallel and Distributed Processing* (2008), IEEE, pp. 1–5.
- [40] MORGENSTERN, M., AND PILZ, H. Useful and useless statistics about viruses and anti-virus programs. In *Proceedings of the CARO Workshop* (2010).
- [41] MURDOCH, S. J. Hot or not: Revealing hidden services by their clock skew. In *Proceedings of the 13th ACM conference on Computer and communications security* (2006), ACM, pp. 27–36.
- [42] NAZARI, A., SEHATBAKHSH, N., ALAM, M., ZAJIC, A., AND PRVULOVIC, M. Eddie: Em-based detection of deviations in program execution. In *Proceedings of the 44th Annual International Symposium on Computer Architecture* (2017), pp. 333–346.
- [43] OR-MEIR, O., NISSIM, N., ELOVICI, Y., AND ROKACH, L. Dynamic malware analysis in the modern era—a state of the art survey. *ACM Comput. Surv.* 52, 5 (Sept. 2019).
- [44] PLATINI, M., ROPARS, T., PELLETIER, B., AND DE PALMA, N. Cpu overheating characterization in hpc systems: a case study. In *2018 IEEE/ACM 8th Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS)* (2018), IEEE, pp. 59–68.
- [45] QUISQUATER, J.-J., AND SAMYDE, D. Electromagnetic analysis (ema): Measures and counter-measures for smart cards. In *Smart Card Programming and Security* (Berlin, Heidelberg, 2001), I. Attali and T. Jensen, Eds., Springer Berlin Heidelberg, pp. 200–210.
- [46] RAVI, S., RAGHUNATHAN, A., KOCHER, P., AND HATTANGADY, S. Security in embedded systems: Design challenges. *ACM Transactions on Embedded Computing Systems (TECS)* 3, 3 (2004), 461–491.
- [47] ROY, N. Owllet: Insect-scale spatial sensing with 3d-printed acoustic structures. *GetMobile: Mobile Computing and Communications* 25, 2 (2021), 14–20.
- [48] ROY, N., AND CHOUDHURY, R. R. Ripple ii: faster communication through physical vibration. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)* (2016), pp. 671–684.
- [49] ROY, N., CHOUDHURY, R. R., AND AL HASSANIEH, H. Causing microphones to detect inaudible sounds and defense against inaudible attacks, June 2 2020. US Patent 10,672,416.
- [50] ROY, N., GOWDA, M., AND CHOUDHURY, R. R. Ripple: Communicating through physical vibration. In *12th USENIX Symposium on Networked Systems Design and*

- Implementation (NSDI 15)* (2015), pp. 265–278.
- [51] ROY, N., HASSANIEH, H., AND CHOUDHURY, R. R. Backdoor: Sounds that a microphone can record, but that humans can't hear. *GetMobile: Mobile Computing and Communications 21*, 4 (2018), 25–29.
- [52] ROY, N., HASSANIEH, H., AND ROY CHOUDHURY, R. Backdoor: Making microphones hear inaudible sounds. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services* (2017), ACM, pp. 2–14.
- [53] ROY, N., HASSANIEH, H., AND ROY CHOUDHURY, R. Riding the non-linearities to record ultrasound with smartphones. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services* (2017), ACM, pp. 189–189.
- [54] ROY, N., AND ROY CHOUDHURY, R. Listening through a vibration motor. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services* (2016), ACM, pp. 57–69.
- [55] ROY, N., SHEN, S., HASSANIEH, H., AND CHOUDHURY, R. R. Inaudible voice commands: The long-range attack and defense. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)* (2018), USENIX Association, pp. 547–560.
- [56] SAMI, S., TAN, S. R. X., DAI, Y., ROY, N., AND HAN, J. Lidarphone: acoustic eavesdropping using a lidar sensor. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems* (2020), pp. 701–702.
- [57] SHETU, R. A., TOHA, T., LUNAR, M. M. R., NURAIN, N., AND AL ISLAM, A. A. Workload-based prediction of cpu temperature and usage for small-scale distributed systems. In *2015 4th International Conference on Computer Science and Network Technology (ICCSNT)* (2015), vol. 1, IEEE, pp. 1090–1093.
- [58] SIVARAMAN, V., CHAN, D., EARL, D., AND BORELLI, R. Smart-phones attacking smart-homes. In *ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)* (2016).
- [59] TAJIK, S., LOHRKE, H., SEIFERT, J.-P., AND BOIT, C. On the power of optical contactless probing: Attacking bitstream encryption of fpgas. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, 2017), CCS '17, Association for Computing Machinery, p. 1661–1674.
- [60] WOOD, D., APHORPE, N., AND FEAMSTER, N. Cleartext data transmissions in consumer IoT medical devices. In *ACM Workshop on Internet of Things Security and Privacy (IoTS&P)* (2017).
- [61] YANG, H., AND TANG, R. Power consumption based android malware detection. *Journal of Electrical and Computer Engineering 2016* (2016).
- [62] YU, T., SEKAR, V., SESHAN, S., AGARWAL, Y., AND XU, C. Handling a trillion (unfixable) flaws on a billion devices: Rethinking network security for the Internet-of-Things. In *Workshop on Hot Topics in Networks (HotNets)* (2015).