# CSAW Embedded Security Challenge Final Team Report: Team Trojan Trackers

Joshua Wang
*Georgia Institute of Technology*
*jwang3453@gatech.edu*

Kayla Kirnon
*Georgia Institute of Technology*
*kkirnon3@gatech.edu*

David Kim
*Georgia Institute of Technology*
*dkim944@gatech.edu*

***Abstract*—This report outlines our team's engagement in the Vertically Integrated Projects (VIP), with a focus on embedded systems security. As part of the course, we participated in the CSAW Embedded Security Challenge, a competition that provided a practical platform to apply the theoretical and technical concepts learned in class. This paper details our approaches to addressing the challenges, the methodologies we employed, and the lessons learned throughout the semester. By linking our work in the competition to the VIP program objectives, we highlight the value of experiential learning in developing advanced problem-solving skills in embedded systems security.**

## Introduction

Our focus within this course was on understanding and mitigating security vulnerabilities in embedded systems. The CSAW Embedded Security Challenge (ESC) served as a capstone, requiring us to apply our knowledge to real-world problems while collaborating as a team. This report details our journey through the competition as part of the VIP, emphasizing our work on theoretical learning and hands-on application.

## Competition Hardware

As the ESC is centered around side-channel attacks in manufacturing systems, the hardware provided closely resembles actual components that could be used in side-channel attacks. We were provided with the following hardware:

- ELEGOO UNO R3 Board ATmega328P with USB Cable
- 12 Voltage 2A 24W Switching Power Supply
- L293D DC Motor Drive Shield Expansion Board
- Stepper motor 12V 350mA
- MG996R Servo Motor 360 control angle
- 120mm DC 12V Computer Fan

Additionally, these parts required some additional setup work once at the lab, which we were able to complete. We were able to use the cyber-physical system comprised of these components to run side-channel attacks using the software provided in `.hex` files on the official ESC GitHub repository. The specifics of each attack for each challenge are detailed in the following sections.

## I. Week 1 Challenges

### A. Normal or Though

*1) Brief Overview:* This challenge features a strange interaction in a New York Pub where an Arduino is given to us

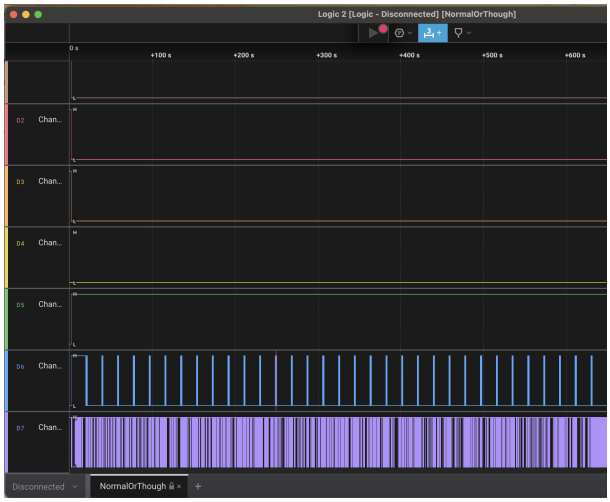


Fig. 1. .sal files within Logic Analyzer

to after our meal, we are left alone in the room to somehow pay, and we are initially confused but then we have some form of realization. The hardware plays a series of sounds that seems to represent whatever payment is happening or whatever happens after we are left alone in the room to pay.

*2) Strategies Used:* Initially, attempting to narrow down the directions to take this in, we tried to think about different methods of payment that the Arduino could have presented once we were left in the room. Before starting to analyze the data, knowing that these were acoustic side channel attacks, we transferred the sound waves into .sal files and utilized saleae's Logic Analyzer software to view information about the sound that was being produced. We also note here that there are two parts to the audio being provided. One is the fan and the other the output from the servo motor. We weren't sure what this would mean later on but made note of it.

1. Frequencies - The first method we tried was using the frequencies in some way to detect if there might have been a possible phone tap or card tap transaction. To this end, we attempted to use the most reoccurring frequencies (49.89 hz was quite popular) in the .sal files to see what devices would normally use those. Unfortunately the range of devices that we were presented with was simply too large to narrow it down, and many of them seemed to be able to do similar things, though there was no way to truly narrow it down, so


Fig. 2. Example usage of cashless payment in a commercial location

we decided to attempt a different route, while keeping this as possible background information.

2. QR Code - when doing research as to what this could possibly be, we realized that some restaurants nowadays have Menus and Payment through QR codes that could lead to online platforms. We decided to look this theory up a little further to see if this was even something that could be created from the audio we received. We attempted to convert the .sal file into a binary file and then to use a python script to generate a qr code from this binary data. Unfortunately up until the point of submitting this paper, the script has been rejecting the data due to its large size not being able to meet the QR code version and size requirements, but we will be continuing to try different methods in hopes that it will work before the competition, as we do believe this is a viable possible solution. We believe that the composition of the binary file we generated with both the motor and fan may be a part of the problem. That or further breaking down the data before attempting to create the QR code.

### B. Friendly Disposition

*1) Brief Overview:* This challenge has us attempt to decode encrypted messages coming from a cyber attack team known as the Malicious Disposition. The messages are meant to follow patterns and we are provided with the first part in this sequence. These parts are described as phases, and we assumed that there are 4 of them by the hint provided to us by the competition leaders.

*2) Strategies Used:* For this challenge, the initial strategy that we used was to record the time values of the sounds from the .sal file we generated and write them out in an attempt to re-classify what the next set shout be. Unfortunately this didn't initially provide anything, and our attempts to continue the sequence were incorrect. The week during which we worked on this was a midterm week before our school's fall break so we decided to reconvene on the challenges the following week, and this was one that stumped us as we moved onto working on the newly released challenge, in hopes of gaining more points there. We did decide that this will be one we attempt to continue working on after submitting this apper, in hopes to have more info for our competition presentation.

## II. Week 2 Challenges

### A. KeyRing 1

*1) Brief Overview:* There are labeled and unlabeled samples. The labeled samples are of 4 different physical keys, and each key has CSV noise sensor data, an audio recording, and an STL file associated with it. There are 40 unlabeled samples, which are either CSV or audio data. The goal is to determine which physical key a given unlabeled sample corresponds to.

*2) Approach and Findings:* As both the audio and sensor recordings contain too much data for a human to feasibly make accurate classifications off of, we chose to use statistical approaches to classify each unlabeled sample.

To preprocess the data and reduce noise, we took sliding windows of data and averaged them. This also helped reduce memory usage to a magnitude feasible for local training. We used a convolutional neural network (CNN) trained on the labeled data to classify the unlabeled data. We found that the audio data only needed roughly 75 epochs to achieve 100% accuracy on the labeled data, whereas the CSV data took roughly 250. Both models used a batch size of 32.

| File | Class | File | Class | File | Class |
|---|---|---|---|---|---|
| 1.csv | KeyA | 21.csv | KeyB | 10.mp3 | KeyA |
| 2.csv | KeyB | 22.csv | KeyD | 11.csv | KeyB |
| 3.mp3 | KeyD | 23.csv | KeyB | 12.csv | KeyA |
| 4.mp3 | KeyC | 24.mp3 | KeyB | 13.mp3 | KeyB |
| 5.mp3 | KeyB | 25.mp3 | KeyB | 14.mp3 | KeyD |
| 6.mp3 | KeyD | 26.mp3 | KeyA | 15.csv | KeyD |
| 7.mp3 | KeyB | 27.csv | KeyB | 16.csv | KeyB |
| 8.mp3 | KeyB | 28.csv | KeyB | 17.mp3 | KeyB |
| 9.csv | KeyD | 29.mp3 | KeyD | 18.csv | KeyB |
| 30.csv | KeyB | 31.csv | KeyD | 19.mp3 | KeyC |
| 32.csv | KeyB | 33.csv | KeyC | 20.mp3 | KeyA |
| 34.csv | KeyB | 35.mp3 | KeyB | 36.mp3 | KeyD |
| 37.mp3 | KeyA | 38.csv | KeyB | 39.csv | KeyB |
| 40.mp3 | KeyC | | | | |

*3) Other Considerations:* We also implemented a Dynamic Time Warp approach to classification using Python and `fastdtw`, and compared the distances from each unlabeled sample to each labeled sample to choose the predicted label as the one with the closest distance. However, this approach did not yield plausible results almost all samples were predicted as Key C.

For CSV recordings in particular, we also considered turning the accelerometer readings across all dimensions into position coordinates by using a double cumulative sum. However, this was not able to recover any meaningful plots, even for the labeled samples.

We noticed that in many of the CSV samples, the values tended to be positive rather than negative, explaining the "line" effect seen in the plot.

## III. Week 3 Challenges

### A. Lizzy

*1) Brief Overview:* We are tasked with decoding an encoded message transmitted by an enemy ship's rotating steel

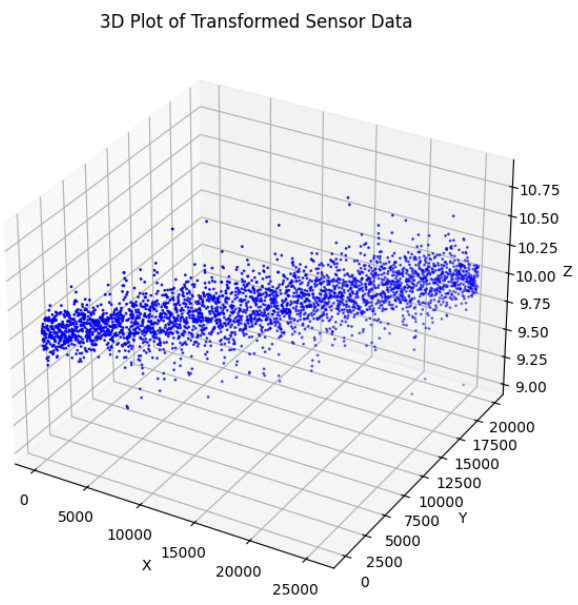


Fig. 3. 3D plot of transformed sensor data

motor contraption. The signal consists of clockwise and counterclockwise rotations, with the key to decoding lying in the duration of these rotations. Observing and interpreting these rotation durations and directions should reveal the message.

*2) Analysis:* We can analyze the rotational directions first. Clockwise and counterclockwise rotations may indicate different values. For example, CW could represent dots or short signals, while CCW could represent dashes or long signals. We can also examine the Duration measurement as the duration of each rotation can be crucial. If the rotations vary in length, it's likely that longer rotations correspond to dashes, and shorter ones correspond to dots. Otherwise, consistent durations might use direction alone to encode the message.

*3) Additional:* Some additional analysis we could perform is to record sequence and duration and then implementation with stopwatch to be more precise. We can also assign code values with longer rotations and shorter rotations or by CW and CWW. We can then translate the message for the outcome.

### B. Conclusion

The enemy signal is encoded through the direction and duration of rotations, likely forming a Morse or binary code. By recording the sequence accurately and interpreting it through Morse or binary translation, we can decode the message. Consistent timing or patterning in rotations will confirm the correct decoding method and reveal the intended message.

### C. Fast Max

*1) Brief Overview:* FastMax aims to unlock two safes protecting a secret chamber within a bank. The first safe requires cracking a 16-digit employee ID encrypted with RSA-like encryption, while the second safe demands a complex alphabetic bank PIN that, when entered correctly, manipulates a combination lock's motor to unlock. Reverse engineering is permitted solely on a FastMaxDummy binary to understand the code logic, but the final solutions rely on exploiting side-channel information, specifically motor acoustics.

*2) Analysis:* Safe 1: Since it is given to us that we will see an RSA-like Encryption, we will have to decode by analyzing the FastMaxDummy hex binary. We can use reverse engineering to identify the encryption scheme to look for public exponent values and modulus size. The binary should reveal if the encryption relies on common RSA values. We can also implement factorization or exploit weaknesses. Depending on the mod size and any patterns identified, factorization may reveal the private key as well. If weak parameters are used, we can use known exploits to speed up the process.
Safe 2: Safe 2 requires acoustic side channel, so the PIN's correctness is verified by motor sounds when the hex file is flashed and ran on the machine. We can find the sound frequency or timing differences to differentiate correct from incorrect characters. Then we can record and analyze motor sounds by using audio recording to capture motor acoustics while entering various character sequences. Analyzing sound duration, frequency, or pauses can hint at correct PIN characters.

*3) Additional:* Additional methods and information we can use are specific tools for reverse engineering like Ghidra and Radare2 to dissect FastMaxDummy hex file to find the RSA encryption information. For audio analysis, Audacity can be used as it helps break down the audio into detailed analysis.

### D. Conclusion

The path to unlocking both safes requires an intricate mix of cryptographic decryption for Safe 1 and acoustic side-channel analysis for Safe 2. Reverse engineering FastMaxDummy hex file will be pivotal to understanding both the RSA parameters for Safe 1 and the acoustic methods for Safe 2. With this combined approach, FastMax could successfully decode the chamber's protective safes and reveal the treasure within.

## IV. Week 4 Challenges

### A. SafeCracker

*1) Brief Overview:* There is an audio file and ino + hex file. The audio file can be analyzed for specific patterns such that the 3-number combination can be determined and replaced in the ino file to output the success message. The provided code in ino file can be run in Ardino IDE to find out what the purpose is and this code includes a variable with placeholder values that correspond to the 3-step combination.

*2) Analysis:* Audio file contains audio that can be analyzed into chunks of waves based on the number of passwords which is 3 and 5 respectively to the 2 parts. We could analyze for peaks which happens 15 times, so It is unlikely to be the case. Therefore, We can analyze the interval or grouping of values like range of numbers and how far apart they are.

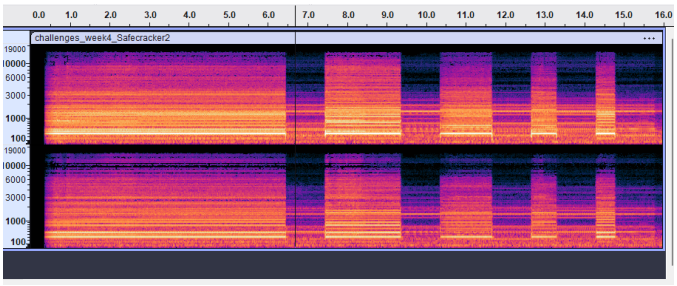


Fig. 4. SafeCracker1 using Audacity

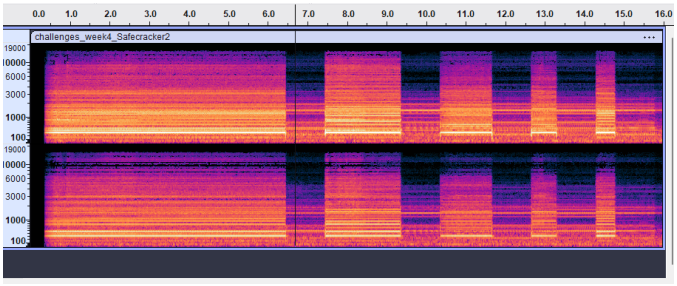


Fig. 5. SafeCracker2 using Audacity

*3) Additional:* We can use Audacity to analyze the waves in depth which is shown in the SafeCracker Figures. This corresponds to a result of the audio for the first and second part of the challenge. Audacity gives the Intervals and Wavelength and all the details of the Audio file in depth. We can also analyze the correlations of each waves and intervals in between.

*4) Conclusion:* This analysis indicates that the 3-number and 5-number combination is embedded within the audio through deeper analysis. These values can act as the key to unlocking the successfully completed message. The step1 and other methods containing the question mark placeholder can be now replaced with the numbers obtained to solve the problem.

### B. Cyber-Physical Systems

A cyber-physical system (CPS) is a system that combines hardware and software to perform a task in the real world. ”Cyber” refers to the software that performs computations and decision-making, and ”physical” refers to the perception of the environment pertinent to the system or the actuation of the decisions made by the software. They have various applications ranging from healthcare to manufacturing, as most of our current world depends on the integration of technology. It is critical to secure cyber-physical systems because they closely interact with their surroundings physically. For example, if a malicious party takes over a cyber-physical system in charge of navigating aircraft, hundreds of lives could be endangered.

### C. Side-Channel Attacks on Cyber-Physical Systems

While Spectre and Meltdown occur at the CPU level and rely on CPU architectural design weaknesses, the scope of SCAs extends far beyond that. Cyber-physical systems pose a promising target for side-channel attacks, as they are much more closely integrated with their surrounding environment. Thus, engineers must consider the additional dimension of easier physical access and observation as part of their threat modeling when designing their CPS. Unfortunately, it is difficult in practice to fully scope out all of the potential attack vectors.

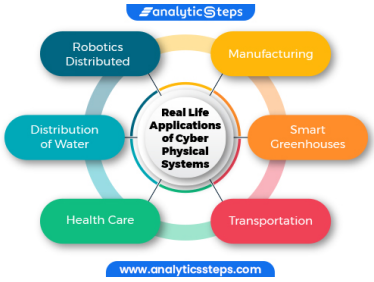


Fig. 6. An introduction to some of the applications of cyber-physical systems in our daily lives. []

This paper will investigate SCAs on cyber-physical systems specifically in manufacturing. This is a critical area for security against SCAs, as CPSs designed for manufacturing are, by nature, designed to interact closely with their surrounding environment and thus pose a considerable attack surface for SCAs.

## V. Attacks and Mitigations

Many types of attacks can be considered side-channel attacks, but some are more commonly discussed and agreed upon. The ones that we will focus on for our paper are acoustic attacks, power analysis attacks, electromagnetic attacks, timing attacks, and network analysis attacks.

### A. Acoustic Attacks

Acoustic attacks take advantage of sound emissions from devices to gain information about the actions being taken by a device. Knowing these actions can enable the hacker to indirectly acquire sensitive information about the user. One example of this is using keyboard sound emissions to calculate the lengths of words typed and the letters that could have been included in words. This method has been used to figure out passwords and decode private messages. Attackers typically use specialized microphones or even smartphones to capture these sound emissions from a distance, applying sophisticated algorithms to analyze the data. The implications of such attacks are particularly concerning in the context of cyber-physical systems (CPS) in manufacturing. Manufacturing equipment often emits distinctive acoustic signatures during operations, such as the sounds of machinery running, robotic arms moving, or tools engaging with materials. An attacker could infer critical information about the production process, operational parameters, and even specific tasks being executed by analyzing these sound emissions. This could lead to unauthorized access to proprietary manufacturing techniques, supply chain disruptions, or even sabotage of operations. Furthermore, as CPS integrates more closely

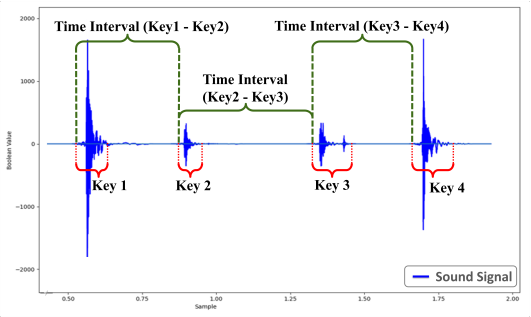


Fig. 7. Acoustic Attacks - An example of acoustic waves that are generated when someone types a 4-letter word on their laptop device. []

with the Internet of Things (IoT), the potential for acoustic attacks increases, posing significant risks to data integrity and manufacturing environments' physical safety.

*1) Mitigations:* An immediately obvious mitigation is to emit additional acoustic noise to prevent relevant acoustic information from being useful to attackers. However, the type of noise and the algorithm used to generate the noise should vary depending on the task at hand. For example, adding a high pitched frequency over a machine that makes low-frequency sounds would not be effective as the high frequencies could be filtered out with a low-pass filter.

Additionally, soundproofing could be installed to reduce the level of acoustic emanations. However, this can be costly and space-inefficient, and the benefits may be limited depending on the threat model determined by the nature of the CPS. In manufacturing, such soundproofing material may not be feasible.

### B. Power Analysis Attacks

Power analysis side-channel attacks rely on the assumption that complementary metal-oxide semiconductors (CMOS) devices have a correlation between the amount of power they consume and what is being done on the device. This type of attack is one of the most powerful attacks performed on hardware. Thus, many researchers around the world are trying to find methods to mitigate their effects by developing resilient hardware. Attackers typically perform these analyses by monitoring the power consumption of a device during its operation, often using oscilloscopes or specialized measurement equipment to capture fine-grained power data over time. However, there are more cost-effective options, that lower the barrier of entry for attackers. There are three main types of power analysis attacks:

*1) Simple Power Analysis (SPA):* This type of power analysis combines prior knowledge of a device's cryptographic algorithm and battery consumption levels over a period of time. This information is then used in mathematical calculations to determine what a device is doing.

*2) Differential Power Analysis (DPA):* Differential power analysis is more useful when a system is more complex than those that would use simple power analysis. This method is best for analyzing sets of measurements to identify data-dependent correlations. It is very effective on tamper-resistant devices, and can handle most systems that simple power analysis can't.

*3) Correlational Power Analysis (CPA):* Correlational Power Analysis is a variation of Differential Power Analysis(DPA). Together, these two types of power attacks have caused much disarray in the world of cybersecurity, as they are capable of breaking encryption algorithms that many once believed could not be broken. Correlational power analysis involves the use of a power model for a given device, which the attacker will then use to find a statistical correlation between the device's power consumption and a predicted output, which will vary based on what the attacker is looking to find or achieve.

*4) Simple Power Analysis Mitigations:* The key to a side-channel attack is the relationship between leaked information and sensitive data. Since simple power analysis involves directly observing the power consumption to identify the pattern, a similar mitigation approach to acoustic attacks can be employed. This involves injecting random noise into the power consumption to obscure the power usage pattern by introducing random delays. Another is to have operations consume a constant amount of power regardless of the input data, making it harder to distinguish between different operations.

*5) Advanced Power Analysis Mitigations:* Differential power analysis is more advanced as an attacker can collect multiple power traces while the device processes different inputs. Similar to simple power analysis, data can be randomized so that operations do not produce the same power consumption patterns, making it harder to find correlations. Another implementation is not correlating certain power consumption with specific operations by balancing power consumption across operations. Finally the Correlational Power analysis takes advantage of correlation coefficients to link the power consumption to specific operations and keys. As with other types of power analysis, the best way to prevent a power analysis attack is to have correlation-immune architecture. Algorithms or cryptography methods should not be easily guessed. Masking techniques can disguise processes, making it difficult for attackers to correlate power consumption with any specific information.

### C. Electromagnetic Attacks

These attacks make use of the electromagnetic radiation that is emitted from devices as they are functioning. In-depth analysis of this data allows attackers to derive confidential information. Electromagnetic attacks can also be conducted without physical access to the device, as they only require a way to capture the emitted electromagnetic signals. This type of attack is highly non-invasive, making it nearly impossible for the user to detect that they have been targeted. Even if they were to become aware, it would be challenging for them to determine what information had been compromised.

*1) Mitigations:* EM attacks can be mitigated by making operations independent or adding clock cycles for all operations.

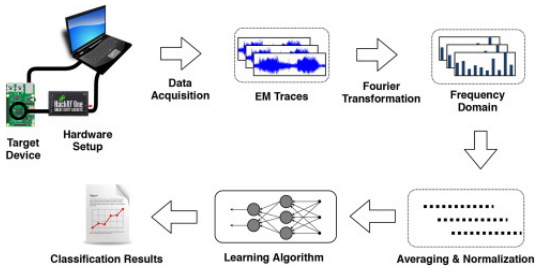


Fig. 8. Electromagnetic Attacks - Some of the possible stages that may be represented when an electromagnetic attack is carried out on a target device. []

This effectively alters the emitted electromagnetic radiation so that it takes an additional step to obtain the information. Adding noise or noise reduction also could further alter the EM signals, making it much harder for the attackers to understand the pattern of the EM signals.

### D. Timing Attacks

Timing attacks are side-channel attacks that exploit the variations in the time it takes for a system to execute cryptographic algorithms. An attacker can learn information about the secret keys being used by measuring the time it takes to perform operations. For example, suppose a decryption process takes longer for certain inputs. In that case, the attacker can deduce which bits of the key might be correct, because it will take longer to get an error message, meaning something inputted might be correct. This vulnerability arises from the fact that implementations often have different execution paths based on input values.

*1) Mitigations:* Timing attacks exploit the time it takes for a system to execute algorithms, and attackers can find the type of algorithm or secret key used based on this information. Therefore, this could be mitigated by setting a time constant for every operation as a delay. This may slow down the overall processing time, but because the pattern is no longer predictable due to spacers of constant delay, timing attacks could be much harder to perform.

### E. Network Analysis Attacks

Network analysis attacks are a type of side channel attack that take advantage of patterns in network traffic to extract sensitive information. By observing details like the timing and size of data packets, attackers can infer what users are doing, what they're communicating, and even cryptographic keys. For instance, if certain operations take longer to process, that timing difference can reveal insights about the data being handled. These attacks highlight how traffic patterns can still leak valuable information even when data is encrypted.

*1) Mitigations:* Network analysis attacks are performed by obtaining patterns in network traffic to extract data. Traffic patterns are vulnerable to leaking valuable information. Therefore, we can add randomized redundant chunks so that the pattern cannot be easily predicted. By doing so, the attacker cannot accurately identify the location with correct sensitive information let alone the pattern. This mitigates significantly the risk of the attacker extracting sensitive data through network traffic.

## VI. Real World Attacks

The usage of attack vectors mentioned in II is highlighted in the disclosure of numerous successful attacks on CPSs used in manufacturing. In this section, we explore two case studies of SCAs in manufacturing and their proposed defenses.

### A. An Optical Attack on Additive Manufacturing Systems

Additive manufacturing (3D printing) produces physical objects from digital models by repeatedly layering material over itself. Consequently, these systems are deeply reliant on software to correctly produce well-formed objects. Software governs the entire process from start to finish – models are designed using computer-aided design (CAD) tools, software determines the best layering technique for printing, and a CPS controls the path and material flow of the *print head*, where raw material is melted and is laid onto previously printed layers.

The attack presented in [**?**] focuses on utilizing optical observations of the print head to recover the printed model, which can be protected intellectual property. To adapt to the often noisy and imperfect observations found in real-world attack scenarios, this attack uses a tailored deep neural network to estimate the position of the print head in three dimensions given optical observations. Using their technique on a baseline print job, the attack was able to digitally recover the printed model to within an average of 0.71 mm of error.

Aside from just presenting an attack, Liang et al. also present several optical defenses to interfere with the recovery of the model. All defenses focused on injecting additional optical noise into the manufacturing process through existing and novel algorithms. For example, the "full power" method of adding optical noise attempts to blind the optical sensors used in the attack to an extent where the neural network degrades in performance.

### B. An Acoustic Attack on Manufacturing Systems

One critical aspect of an SCA's feasibility is the ease of collecting the information necessary for the SCA. If the SCA exists but requires complex sensors and observational instruments to carry it out, then its practical impact is reduced. On the other hand, an easy-to-perform attack is an extremely dangerous one. This section will review an attack on manufacturing systems that can be carried out using only a cell phone and its sensors [**?**]. Due to the widespread usage of phones at factories, the authors of this attack argue that an attacker could spread malicious software to enable phones to listen to and transmit factory floor information.

At its core, the attack relies on acoustic and magnetic sensors embedded in a phone to record the operation of manufacturing systems and transmit them over a phone call. Then, elementary pattern-matching and more advanced machine-learning algorithms are used to recover the manufacturing jobs.

Unlike the attack by Liang et al. on additive manufacturing systems, Hojjati et al. demonstrate that their attack also performs well on CNC mills and argue that it can generalize to other forms of manufacturing CPSs.

To mitigate against these attacks, the authors recommend playing pre-recorded manufacturing noises during manufacturing. This is because noise reduction algorithms can circumvent the tactic of playing randomly generated noise.

### C. Do imperfect mitigations make a difference?

A common mitigation to SCAs in manufacturing CPSs is to add some sort of noise to degrade the accuracy of the techniques used to recover the manufacturing job. However, there is a clear trade-off between the quality of mitigations used and the price of running a single manufacturing job. To balance this trade-off effectively, it is important to keep the overall goal of an attacker in mind: to leak the manufacturing *details* that could not been known otherwise (such as through observation of the final product). To that extent, adding in enough noise so that attackers cannot recover the fine details of a manufacturing job should be, at minimum, sufficient as a mitigation.

## VII. Group Achievements

We have accomplished several key milestones in preparation and participation of the CSAW Embedded Security Challenge (ESC) this year. First, we conducted a thorough review into side-channel attacks and possible approaches and mitigations in the context of cyber-physical systems used in manufacturing. Using this information, we wrote and submitted a detailed report containing our findings and additional commentary as a CSAW ESC qualification paper. As a result of our thorough investigation, we were accepted to proceed to the final stage of the ESC.

While writing the qualification report, we also familiarized ourselves with challenges from the final stages of ESC in previous years. Among all members, all available challenges from ESC 2023 (the previous year) were reviewed. Although the techniques used previously are not directly applicable to this year's challenge of only using side-channel attacks, the insights gained through this process are immensely valuable to both our success in the ESC final as well as the development of our personal skills.

The final stage of the ESC involves performing real side-channel attacks on a cyber-physical system comprised of an Arduino Uno and other components. To prepare for the final stage, we researched possible configurations and attacks that could be tested using the combinations of the components, and presented our findings to the CSAW teams. We found numerous resources online detailing example side-channel attacks carried out on Arduino Unos, including a paper which applied neural networks to attack the exact chip that was mentioned in this year's challenge components.

During the Embedded Security Challenge, we applied our knowledge to conduct side-channel attacks on provided hardware. First, we set up the hardware manually by stripping wires and fastening wire holders. Then, we were able to approach the rest of the challenges, which were gradually released over the span of four weeks.

The final stage of the ESC also includes presenting to judges at the annual CSAW conference in New York, as well as a poster session where we would present to other teams. In order to prepare for this, we created a detailed slideshow and poster containing all of our insights into this year's challenges.

At CSAW, Kayla and David represented the team in front of a panel of judges. Though we did not win, the judges informed us that one of our challenge solutions had the most accurate solution in the US region.

## VIII. Individual Contributions

### A. Joshua Wang

Joshua successfully led our team of three to compete in the 2025 CSAW ESC qualifications and finals. He led team meetings and made sure members were completing tasks on time and satisfactorily. He also handled logistics such as submitting the paper on the portal and communicating with the project mentor and other project teams. Additionally, he helped create slides for the midterm presentation which detailed the team's progress so far.

When writing the paper, he created the outline and had several suggestions for paper flow in order to create a compelling and well-written report. He conducted background research on the history of side-channel attacks, including non-manufacturing related side-channels such as Meltdown and Spectre, in order to complete the Introduction section. To highlight the importance of side-channel attacks in manufacturing CPSs, he found two case studies and conducted a detailed analysis of both of them. Finally, he additionally wrote the conclusion, abstract, and parts of the mitigations sections in the submitted paper.

During and after the paper submission period, he also investigated previous challenges from ESC 2023 using resources available in the VIP GitHub repository. Investigating these challenges gave him a good sense of what the second stage of the competition would be like.

Joshua conducted OSINT on the target devices at the request of our mentor. He found numerous resources online detailing how to conduct side-channel attacks on Arduinos, and even one on the specific chip using neural networks. Additionally, he came up with hypothetical attacks using the other components presented in the shipping list given to us, such as using visual observations of the motors while underpowered to observe when expensive CPU computations were being performed.

For the final stages of the competition, he led the rest of the team to complete the challenges given and present the team's work to CSAW judges at NYU. He set up the hardware in the lab, and also worked on the machine-learning related challenges in week 2. His solution using convolutional neural networks was the most accurate in the US region, and he also presented 2 other approaches that were less effective. Additionally, he also helped create the poster, presentation,

and write portions of the final competition paper submitted to judges. Finally, he created portions of the final class presentation and wrote and edited portions of the final class paper.

### B. Kayla Kirnon

Kayla's time in this VIP semester has been split three ways, between communicating and working with the teammates, working on deliverables such as the qualification report, the first team presentation, the midterm report, the team's slide deck for the final competition, the team's poster for the finals, and now this final paper, as well as individual research, to ensure she had the knowledge to assist her team adequately throughout the course of the competition and semester.

Team communication outside of our meeting times has mostly been through Slack but also via imessage, once we realized we would all be more responsive on that platform. Our team meets every week on Thursday and she has attended all of these meetings for the semester. This is where she was able to help choose the team name, provide my teammates with information about my previous experiences for the competition, and get updates from our team leader on what she should be working on in a given week. She also was able to co-ordinate the work needed to fully prepare us for the finals trip in terms of the slides and poster creation. She has also attended all of the broader CSAW meetings, as well as the VIP full class meetings.

In each of the deliverables we have worked on so far, Kayla had specific tasks which I have completed. For the first class presentation, she created the cover slide, the overview slide, as well as the slides containing my personal updates. For the qualification paper, and this midterm paper, she provided all of the in-depth definitions and information for each of the 5 side channel attacks that we decided to cover. She also made sure to adapt the paragraphs to relate to manufacturing systems in keeping with the competition theme, and updated the abstract and paper introduction.

In terms of research, since she competed in last year's competition on side channel attacks, She initially began researching attacks that included multiple different types of side channel attacks, but she realized that she needed a refresher on the various types to be able to better identify them. Kayla then pivoted to researching each type of attack individually and that helped with writing my section of these papers and better prepared me to identify what each of the parts of the Arduino Uno being sent to us could be used for. She is excited for the hardware to arrive and the challenges to be released and is looking forward to working even more closely with my teammates to solve them.

### C. David Kim

As a part of the Embedded Systems VIP team, he has logged the team's progress and list of things to do in our GitHub page through instructions given in slack. As a group, the team has completed research on Side Channel Attacks and its mitigation and has had weekly meetings to keep note of our work. The meeting times at first have been in question due to overlapping classes for some students, but it has been quickly settled.

First, to detail the work that individual members have been assigned, the team was to first learn about Embedded Systems and its importance. This was crucial as some members had limited knowledge about this topic coming in. David did not have any background knowledge and came in expecting to learn new things, so he believes defining the topic was a very good way to start this VIP.

The team then worked on Qualification papers which dealt with Side Channel Attacks that are exploited usually in embedded systems. David was specifically tasked with the mitigation part of this. He first learned about the different types of Side Channel Attacks that exist and researched those specifically to see if there were any flaws or methods that could be utilized to mitigate the attack.

As the team was finishing up, they were assigned presentation for the progress, which went smooth for all. Just like the qualification paper, David dealt with Side Channel Attacks and its mitigation in the slide along with his personal insights and progress.

As we transitioned to the CSAW Embedded Security Challenge, I took on a more hands-on role during weeks 3 and 4 of the competition. These weeks posed complex challenges, including week 3 which involved analyzing motor rotation patterns to extract encoded messages and week 4 with analyzing audio to uncover a secret message. My efforts here focused on leveraging audio analysis tools and debugging strategies to identify and decode the necessary patterns efficiently. The team created a presentation and a poster to be used for presentation at the competition.

Returning from the CSAW competition, our team was proud of the progress we made, even though we encountered challenges that tested our problem-solving abilities. Each challenge reinforced concepts we had explored during the VIP course, offering practical applications and deepening our understanding of embedded systems security.

Finally, David tracked all my progress and things to do in the GitHub repository. This helped him organize things into something he can understand at a glance. For the Notebook, he logged his progress for each week and updated Github Issues weekly based on tasks. The team was happy to be qualified for the finalist round and they will be going in person to take part in the competition. The team is looking forward to learning more about embedded systems and getting experience from the competition.

## IX. Conclusion

Participating in the CSAW Embedded Security Challenge was a pivotal component of our work in the Vertically Integrated Projects (VIP) program. This competition served as a practical application of the concepts and skills we have developed throughout the course, reinforcing the critical connection between classroom theory and real-world embedded systems security challenges.

The competition provided a structured opportunity to explore advanced topics such as side-channel attacks, cryptographic analysis, and hardware reverse engineering. It allowed us to tackle realistic security problems, applying tools and techniques that are directly relevant to the field. As part of the VIP course, the competition emphasized collaborative problem-solving and iterative development, skills that are fundamental to success in both academic and professional environments.

While the challenges tested our knowledge and creativity, they also revealed areas for growth, particularly in time management and refining our approaches to complex problems. These lessons will inform our future contributions to the VIP program as we continue to build on the skills and insights gained from this experience.

Overall, our participation in the CSAW competition not only fulfilled a core requirement of the VIP course but also solidified our interest and expertise in embedded systems security. We are grateful for the support and guidance of the VIP faculty, and we look forward to applying these experiences in the remainder of the program and beyond.

## ACKNOWLEDGMENT

We'd like to extend a special thank you to our team advisor Kevin Thompson of Georgia Tech Research Institute for answering our questions and assisting when needed, as well as talking us through multiple different scenarios.

## REFERENCES

[1] QrCode Tiger, QR Codes for Payments, https://www.qrcode-tiger.com/qr-code-for-payment.

[2] "How are QR codes generated?", https://www.qr-code-generator.com/blog/how-are-qr-codes-generated/

[3] "Generating QR codes in python", https://realpython.com/python-generate-qr-code/.