

DeepAuditor: Distributed Online Intrusion Detection System for IoT Devices via Power Side-channel Auditing

Woosub Jung
William & Mary
wsjung@cs.wm.edu

Yizhou Feng
Old Dominion University
yfeng002@odu.edu

Sabbir A Khan
Old Dominion University
skhan@cs.odu.edu

Chunsheng Xin
Old Dominion University
cxin@odu.edu

Danella Zhao
Old Dominion University
zhao@cs.odu.edu

Gang Zhou
William & Mary
gzhou@cs.wm.edu

ABSTRACT

As the number of IoT devices has increased rapidly, IoT botnets have exploited the vulnerabilities of IoT devices. However, it is still challenging to detect the initial intrusion on IoT devices prior to massive attacks. Recent studies have utilized power side-channel information to identify this intrusion behavior on IoT devices but still lack accurate models in real-time for ubiquitous botnet detection.

We propose the first online intrusion detection system called DeepAuditor for multiple IoT devices via power auditing. To develop the real-time system, we propose a lightweight power auditing device called Power Auditor. We also design a distributed CNN classifier for online inference in a laboratory setting. In order to protect data leakage and reduce networking redundancy, we then propose a privacy-preserved inference protocol via Packed Homomorphic Encryption and a sliding window protocol in our system. The classification accuracy and processing time are measured, and the proposed classifier outperforms a baseline classifier, especially against unseen patterns. We also demonstrate that the distributed CNN design is secure against any distributed components. Overall, the measurements are shown to the feasibility of our real-time distributed system for intrusion detection on IoT devices.

KEYWORDS

Power Intrusion Detection System, Distributed Online System, Privacy-preservation, Internet of Things

1 INTRODUCTION

Internet of Things (IoT) devices have become the new cybercrime intermediaries between attackers and users to process cyberattacks. For example, in October 2016, a massive distributed denial-of-service (DDoS) attack incapacitated the Domain Name System provider Dyn. This made several Internet platforms and services, such as Amazon, Netflix, PayPal, and Twitter, temporarily unreachable to numerous users in Europe and North America. This IoT botnet attack was called Mirai and exceeded 600 Gbps in volume at its peak. This overwhelming amount of the traffic was sourced from 65,000 injected IoT devices, including routers, security cameras, and digital video recorders [6]. These IoT devices were known at the time to have weak security protection and to be vulnerable to attacks. As reported by Symantec [47], thousands of outdated routers were targeted by the worms exploiting their vulnerability. Since then, many variants of Mirai have emerged to target the weaknesses of IoT devices. Besides serving as the intermediaries of

DDoS attacks, IoT devices were also found to serve as attack proxies for multiple cybercrimes, such as clickjacking and spearphishing [41][40].

Even though the Mirai attack happened five years before the time of writing, IoT devices are still vulnerable to IoT botnets. Mirai and its variants used a simple brute-force attack [46] when first transforming IoT devices into their bots, which is still cybercriminals' preferred option for executing attacks [28]. Despite the clear intrusion procedures of botnet attacks on IoT devices, it is not easy to determine whether an intrusion has occurred. The main concern is that the network traffic generated on endpoint devices is not noticeable as malicious behavior in the initial stages of the attack. In this situation, deploying network-based botnet detection systems into different IoT devices/vendors is heavyweight and intrusive to users. For example, 84 different IoT devices/vendors were found to engage in the Mirai bots. Moreover, they were related to more than 300 different communication protocols and platforms [6]. Thus, network-based botnet detection requires a great deal of modification in programming languages or operating systems on diverse IoT devices. One approach to tackling this issue involves the use of power side-channel information for the detection of stealthy attacks. Using a power side-channel is durable and universal since power traces are hard to compromise and can capture accumulated tasks on heterogeneous devices, such as different hardware, vendors, operating systems, etc. Meanwhile, it is nearly impossible for adversaries to mimic normal power draw behavior while attacks. Thus, some pioneering work in this area utilized power side-channel data to detect malignant behavior on mobile devices in the early 2010s [27][51]. However, these studies are outdated and need to be validated in IoT environments.

Several recent studies have utilized power auditing — the analysis of power consumption — to explore IoT security and defend against malicious attacks [36][29]. However, those proposed systems mainly focused on detecting massive DDoS attacks on IoT devices. It is still therefore challenging to distinguish the subtle initial stages of an IoT botnet intrusion. Jung et al. [25] introduced IoT botnet detection via power modeling. In this research, the authors designed a Convolutional Neural Network (CNN) model using one-dimensional power side-channel data to classify malicious intrusion. While the CNN classifier showed promise in detecting subtle differences in power traces, the fundamental problem with this study is that it was conducted offline with a bulky and expensive power monitor. Thus, it is not practical for ubiquitous botnet detection on IoT devices. To tackle these challenges, we designed a lightweight

power auditing device and a distributed online CNN classification system for resource-constrained IoT devices. Our proposed end-to-end system *DeepAuditor* was developed in a distributed setting and can simultaneously detect the initial botnet intrusion on multiple IoT devices via power auditing.

With this aim, our research questions in this paper are as follows:

- How can we audit power side-channel information of IoT devices in real-time for ubiquitous botnet detection?
- How can we conduct online inference for multiple IoT devices in a distributed setting?
- How can the distributed classifier prevent data leakage and networking redundancy?

To answer the first question, we designed a small form-factor device called Power Auditor that is capable of measuring the power traces of a connected IoT device for behavior classification. Unlike off-the-shelf power monitors, our Power Auditor is lightweight and portable. Thus, we envision devices such as Power Auditor as an important component of future smart plugs. Nowadays, smart plugs are often used in IoT environments to control electronics remotely for the sake of convenience, e.g., Amazon Smart Plug [15]. As IoT devices are getting popular, the smart plug market also grows quickly; its compounding annual growth rate is approximately 42% over the forecast period between 2021 and 2026 [2]. According to our prototype performance, the proposed algorithms are lightweight and therefore can easily be integrated into future smart plugs for ubiquitous botnet detection.

To address the other research questions, we developed distributed CNN classifier components: *Data Inferencer* in a user site and *Computing Cloud* in a cloud site. The proposed CNN design outperforms the baseline classifier; we increased the classification accuracy by up to 17% in leave-one-out tests. In our leave-one-device-out tests, a classifier is trained with a certain device-type dataset and tested with the remaining device-type data to show its robustness against a new device. The leave-one-botnet-out tests are also crucial concerning practical deployments because the classifier has to be robust against unseen attack patterns but not overfitted. We then designed distributed protocols between our CNN components: a privacy-preserved CNN protocol and a sliding window protocol. Note that side-channel information can also reveal users' private data to the cloud site [31]. To remedy this concern, we designed the privacy-preserved CNN protocol in our distributed system. This protocol also addresses the problem of attackers extracting classifier model parameters from distributed environments [44] [43]. The sliding window protocol was developed to reduce networking redundancy.

In summary, our contributions in this study are threefold.

- As a proof of concept, we designed a dedicated small form-factor device called Power Auditor to realize online botnet detection. The Power Auditor is lightweight and portable, and therefore can easily be integrated into future smart plugs.
- We are the first to develop a distributed online intrusion detection system for IoT devices via power auditing. We designed distributed CNN components and proposed distributed protocols between user and cloud sites in order to conduct real-time inference without data leakage.
- We demonstrated the performance of our classification system in a laboratory setting: 1) The proposed CNN classifier

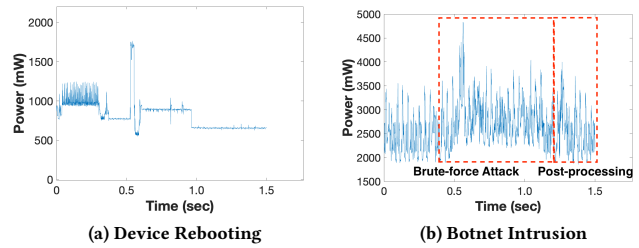


Figure 1: Power Traces collected by our Power Auditor

detects malicious behavior with an accuracy of up to 98.9%, which outperforms the baseline, 2) we theoretically analyzed the data protection of the privacy-preserved protocol, and 3) the distributed system supports about one hundred IoT devices simultaneously by using our laboratory server.

The rest of the paper is organized as follows: Section 2 provides background and a threat model for our study. In Section 3, we introduce our botnet detection system *DeepAuditor*. In Section 4, we present the Power Auditor design of our system. Section 5 introduces the distributed online CNN model for botnet detection. In Section 6, we describe online system implementation. Section 7 demonstrates online performance evaluation of our system. Section 8 presents our thoughts regarding limitations and future work. Section 9 summarizes related work. Finally, we conclude this paper in Section 10.

2 BACKGROUND AND THREAT MODEL

This section provides background knowledge and introduces a threat model.

2.1 Intrusion Detection via Power Modeling

Power traces can capture accumulated tasks to identify abnormal behavior. Several recent studies demonstrated that subtle activities on smartphones can be inferred from power consumption data [27, 50, 51]. Likewise, other studies on IoT security [25][29] explored how IoT devices' behavior generates different patterns of power consumption data. For example, Figure 1 illustrates two examples of different activities collected by our Power Auditor. As shown in Figure 1a, rebooting IoT devices generates power traces that are distinct from the power patterns that occur in Mirai botnet intrusion [Figure 1b]. In Figure 1b, an attacking bot enters the device using different username/password combinations and conducts post-processing. This generates two power traces: one from the brute-force attack and one from the post-processing. Furthermore, since the Mirai botnet family utilizes brute force attacks [48], the power traces generated by these intrusions look nearly identical across dataset. Thus, it is feasible to identify malign behavior by analyzing power traces on IoT devices.

In IoT botnet detection, it is crucial to detect this intrusion behavior in its early stages. Otherwise, botnets grow so quickly, so it will be too late to defend against DDoS attacks. We aim to identify whether the real-time power traces suggest malicious or benign behavior and classify them accordingly. Though RNN could also help in streaming data, it is more suitable in natural language applications. Instead, CNN is widely used in sensor streaming applications and classifies instances very fast in the inference stages. Therefore,

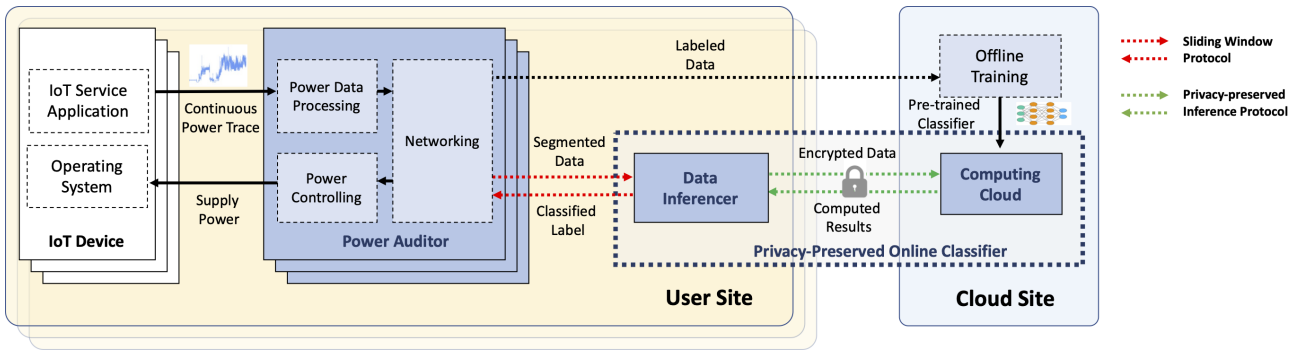


Figure 2: DeepAuditor System Overview – The system consists of three subsystems with two distributed protocols.

we designed a CNN design to realize an online botnet intrusion detection system. More detailed design factors will be discussed in the following sections.

2.2 Threat Model

2.2.1 Client-side Vulnerabilities. In DeepAuditor, we train power traces generated by existing IoT botnets. However, we also assume that an adversary is capable of conducting various patterns of botnet attacks. Thus, we consider two possible strategies that the adversary can use to attack our client-side. 1) Exploit vulnerabilities in the Power Auditor device. 2) Generate complicated post-processing jobs that create unseen power patterns. For example, downloading unexpected files or connecting to multiple servers can create more complicated data patterns.

To minimize the vulnerabilities of the Power Auditor, our proposed Power Auditor only monitors the power consumption of the physically connected IoT device within a local network; thus, the Power Auditor does not allow any inbound traffic from remote sources. This assumption is especially valid when smart plugs also do not allow users to access ssh/telnet services [2]. Instead, these smart plugs are mostly managed by manufacturer apps. As we envision Power Auditor to be an integral part of future smart plugs, both Power Auditor and the smart plugs are secure against brute-force attacks. To address the post-processing job vulnerability, segmented data from different patterns were trained as botnet instances in our deep learning model. Thus, as long as power side-channel information is noticeable enough to label, our system is capable of detecting a diverse set of botnets beyond that are well-known. Our leave-one-out tests also demonstrate the classifier performance against unseen patterns in Section 5.1.

2.2.2 Server-side Vulnerabilities. In addition to exploiting the above vulnerabilities, adversaries could also target our cloud servers. Note that we implemented our classification model into two cloud-based edges. Thus, we assume any user-held application and model-held server in our system can become a semi-honest adversary. This means that they may try to steal information from received messages. For example, servers may infer IoT device behavior based on the power trace input, and users try to learn the server’s model parameters based on the server output. We consider all parties non-colluding for their input data and output data. It is essential for our system to avoid user’s privacy data disclosure that leads to poor credibility.

The emerging attacks presented in [43, 44] are also threats that we need to consider in our model. User-side can launch the model extraction attack [44] to extract the convolution layer and fully connected parameters based on the server received message. Server can process membership inferences attack [43] to compare the user input with the server’s pre-trained dataset. In this research, our privacy-preserving mechanism masks the intermediate/final output for both user and server. However, the user still can learn the correct predicted result. Simultaneously, our privacy-preserving mechanism protects the server holds model parameters from the user, and user input is oblivious for server. We applied a flexible method to protect the output correctness and prove our system security by using a real-ideal paradigm [38], as introduced in detail in Section 7.3.

3 DEEP AUDITOR SYSTEM OVERVIEW

In this section, we introduce the distributed IoT botnet detection system DeepAuditor. As shown in Figure 2, the proposed system consists of three subsystems: Power Auditor, Data Inferencer, and Computing Cloud. The ultimate goal of DeepAuditor is to identify subtle power differences in real-time between normal behavior and IoT botnet intrusions on multiple IoT devices. To build the system, Power Auditors are used to secure power-trace data of IoT devices. Data Inferencer and Computing Cloud are then used as online classifiers that can process intrusion detection. Altogether, the proposed distributed components accomplish the online intrusion detection via power side-channel auditing.

In the user site, the Power Auditor is connected to each IoT device to collect power consumption traces; there can be multiple Power Auditors for the corresponding IoT devices. We propose several functionalities of the Power Auditor for intrusion detection as follows. First, the Power Auditor is capable of auditing a device’s power consumption footprint via the Power Data Processing module. During the online phase, the Networking module then transmits the collected data to the Data Inferencer for online classification. Finally, the Power Controlling module supplies power to the connected IoT device. Section 4.1 introduces the universal design of the Power Auditor in detail.

Next, we developed a sliding window protocol between the Power Auditor and the Data Inferencer in the user site. This is because overlapping input instances of the CNN classifier can help to achieve better detection accuracy. However, if clients send those

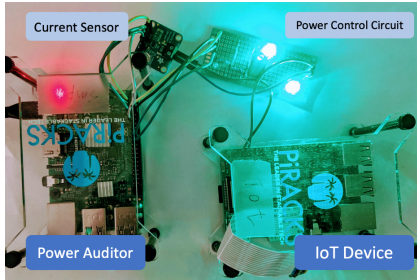


Figure 3: Power Auditor Prototype

overlapping windows to the server sides, this will lead to networking redundancy. For example, if a Power Auditor sends 1.5 seconds of data every 0.5 seconds, this will bring 67% networking redundancy. Instead, Power Auditors send segmented data to the Data Inferencer that concatenates the segmented packets for assembling input instances. Overall, in order to meet the real-time classification, our DeepAuditor needs to process each input instance less than the size of the sliding window. More details will be described in Section 4.2.

We then designed a one-dimensional CNN classifier for botnet intrusion detection, which we explain in Section 5.1. Our evaluation results demonstrate that our classifier outperforms the state-of-the-art classifier [25] for power modeling. Based on the proposed 1-D CNN model, we implemented and deployed the pre-trained CNN classifier into a distributed lab setting.

Lastly, Section 5.2 introduces the Privacy-Preserved CNN inference protocol for online prediction in a distributed setting. This protocol enables the CNN communications between user and cloud sites without leaking data. The Data Inferencer runs in the user site, receives power traces from Power Auditors, and sends encrypted data to the Computing Cloud to offload computation resources. In the cloud site, the Computing Cloud is in charge of CNN inference computations, i.e., nonlinear activation computations in the CNN. Then, the Computing Cloud sends the encrypted computation results to the Data Inferencer for online inference. Finally, the Data Inferencer classifies the final prediction of whether the given input power trace is benign or malicious.

4 POWER AUDITOR DESIGN

In section 4.1, we introduce the three universal software components of the Power Auditor: Power Controlling, Power Data Processing, and Networking modules. In section 4.2, we then describe the data transmission protocol between the Power Auditor and the Privacy-Preserved CNN Classifier.

4.1 Component Design

In our Power Monitor, the Power Controlling module supplies power to a connected IoT device. Next, the Power Data Processing module reads power consumption traces of the attached IoT device. The Networking module communicates with the server-side to convey the sensing data for online classification. Figure 3 shows our prototype of a Power Auditor and a connected IoT device. All the source codes for Power Auditors will be available for download.

The Power Auditor bypasses power to the connected IoT device from an AC adapter. As shown in Figure 4, we built a FET-based

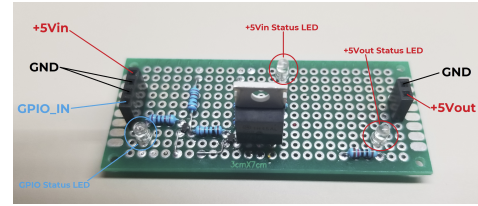
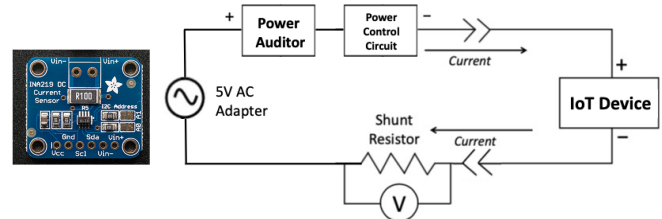


Figure 4: Power Control Circuit



(a) INA219 Sensor

(b) Logical Circuit of Power Measurement

Figure 5: Power Measurement Design

switch to turn on the connected IoT device via GPIO pin from the Power Auditor device. The current draw of this module is only 20mA when the power output is being provided, while a Power Auditor supplies enough power to most IoT devices. Therefore, IoT devices do not need any modifications in our system.

We also designed a current circuit for power measurement. To get the current and voltage reading on the connected IoT device, we utilized the current sensor INA219 [5], as displayed in Figure 5a. This sensor includes a shunt resistor and provides ADC conversion to the Power Auditor. Figure 5b illustrates how the current sensor provides power consumption data of the IoT device. In this circuit, Power Auditor measures the voltage drop around the shunt resistor at a high frequency. Based on this data, we calculate the current values going through the entire circuit. By doing so, we can measure the power consumption of the IoT device since the voltage input is also fixed. According to the specification of the sensor used, the maximum error rate is approximately 0.5%, which is negligible.

After measuring power data, the Power Data Processing module pushes each power reading instance into a local queue. Then, the Networking module fetches the queued data periodically and assembles the collected data to a TCP packet for data transmission to the online classifier. Detailed networking protocol and interface format are illustrated in the next subsection.

4.2 Sliding Window Protocol

We introduce the interface design between the Power Auditor and the Data Inferencer for online inference. First, we determined a window size of 1.5 seconds for botnet intrusion input. Jung et al. [25] revealed that Mirai and its variants have similar time distributions for the initial intrusion during the propagation period, which is less than 1.5 seconds. It should be noted that this invasion time may vary depending on systems or botnets. However, as long as it is noticeable to label, a window size would not be a critical issue.

Next, we applied a sliding window with one-third overlap for better classification accuracy. Our reasoning is that malicious instances can be truncated within a single window. Different overlapping

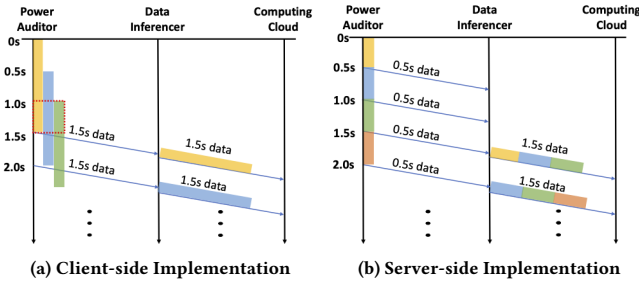


Figure 6: Sliding Window Protocol Options— We adopted the Server-side scheme to reduce networking redundancy

```

smpg1,4da77a50-aeaf-11,1700,1500,3,850,1650.000,1864.024,2380.793,1819.207,1736.890,
1746.951,1675.610,1755.183,1765.244, ..., 2574.695,2565.549
    
```

Figure 7: TCP Packet Example in Sliding Window Scheme

ratios are also possible but one needs to consider the trade-off between classification accuracy and network bandwidth. We were able to achieve real-time prediction with the one-third overlapping ratio. Figure 6 illustrates the overlapping sliding window scheme in our DeepAuditor system. As shown in Figure 6a, if a Power Auditor transmits a data packet of 1.5 seconds every 0.5 seconds, this will create redundant packets. Instead, a Power Auditor reads 0.5 seconds of data and sends it out once collected, as illustrated in Figure 6b. The Data Inferencer then receives the packet every 0.5 seconds. After receiving three consecutive packets, the Data Inferencer assembles the last three packets and feeds them into the pre-trained CNN classifier for online inference. By doing so, we avoided unnecessary network redundancy in a distributed setting.

We then implemented the interface format accordingly. Table 1 illustrates the packet header format. The number of data points is determined based on the following values. For example, our sampling rate is 1700, and the window size is 1.5 seconds. Consequently, the number of data points in a single instance for classification is $1700 \times 1.5 = 2550$. Since we adopted the server-side sliding window scheme, the Power Auditor also set the Sliding Window Ratio header to 3. Finally, the number of data points in a single TCP packet is $2550 \div 3 = 850$ in our system. The message body contains a list of the power-sensing data. Figure 7 describes an example of the raw TCP packet data.

5 DISTRIBUTED CNN CLASSIFIER DESIGN

In Section 5.1, we present a one-dimensional CNN architecture. Section 5.2 then describes a distributed inference protocol between the DeepAuditor components that offloads CNN computations and protects data leakage.

Table 1: TCP Packet Interface Format

Header	Type	Description	Example
Hostname	String	Hostname of Power Auditor	smpg1
Message ID	String	Unique ID for each message	4da77a50-aeaf-11
Sampling Rate	Integer	Sampling Rate of Current Sensor (Hz)	1700 (Hz)
Window Size	Integer	Length of a single window (ms)	1500 (ms)
Sliding Window Ratio	Integer	Ratio of Overlapping Data	1: No Overlapping 2: 1/2 Overlapping 3: 1/3 Overlapping
Number of Data Points	Integer	The number of data points in a single TCP packet	$850 (\text{Sampling Rate} \times \text{Window Size} / 1000 / \text{Sliding Window Ratio})$
Data Points	Float	List of Power consumption data (mW)	1854.878, ... (mW)

5.1 1-D CNN Classifier Design

We designed a 1-D Convolutional Neural Network (CNN) architecture for botnet intrusion detection that outperforms the state-of-the-art model (CHASE'19) [25]. Jung et al. introduced a CNN model and conducted an offline evaluation. Although this CNN classifier showed promise in identifying subtle differences in power traces, this model performed poorly in some experiments, especially against unseen patterns. For example, the leave-one-out tests introduced a prediction accuracy of below 80%, which suggests a possible overfitting problem. Indeed, achieving high accuracy in the leave-one-out tests is very important in our system. In a real-world deployment, it is not possible to train all the cases or botnets beforehand. Therefore, the poor leave-one-out results motivated us to propose an enhanced CNN model to avoid overfitting and increase classification accuracy for IoT botnet intrusion.

5.1.1 CNN Design Philosophy. Figure 8 provides an overview of the proposed 1-D CNN architecture. This CNN classifies time-series power-trace input as belonging to one of four behavior classes in IoT devices: IoT Service, Idle, Reboot, Botnet Intrusion. Thus, we seek to distinguish malicious intrusions from other common behavior.

Because the DeepAuditor processes power consumption data, the input layer prepares one-dimensional input to feed into the convolutional layer. The power sensing model has a sampling rate of 1.7kHz and a window length of 1.5 seconds (See Section 4.2). Thus, the input instance size is (1×2550) . We then use 32 one-dimensional (1×128) kernels at a stride size of 64. Consequently, the convolutional layer computes a dot product between the power traces and 32 kernels.

Moving forward, we carefully chose the design factors of the 1-D CNN to solve the complex problem of botnet intrusion detection while avoiding overfitting. To demonstrate the effectiveness of each design factor, we used a public dataset, which was collected in 2019, for botnet intrusion detection [1] containing power trace data from four different behavior classes (IoT Service, Idle, Reboot, Botnet Intrusion) from three types of IoT devices (Security Camera, Router, Voice Assistant) and three IoT botnets (Mirai, Sora, Masuta). In the dataset, each class has 2,000 instances of 1.5 seconds of power-consumption data. The CHASE'19 CNN classifier performs well in self-evaluation tests where the training and test data are from the same type of device. However, this CNN architecture performs poorly in leave-one-out tests where the test dataset is from a different type of device than the training dataset. We can infer from this that the CNN model is either not powerful enough or overfitted to the specific datasets.

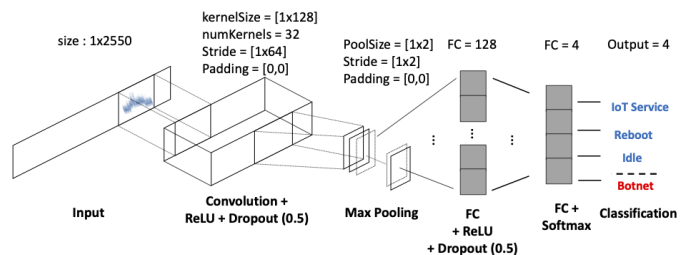


Figure 8: 1-D CNN Model for the power-trace classification

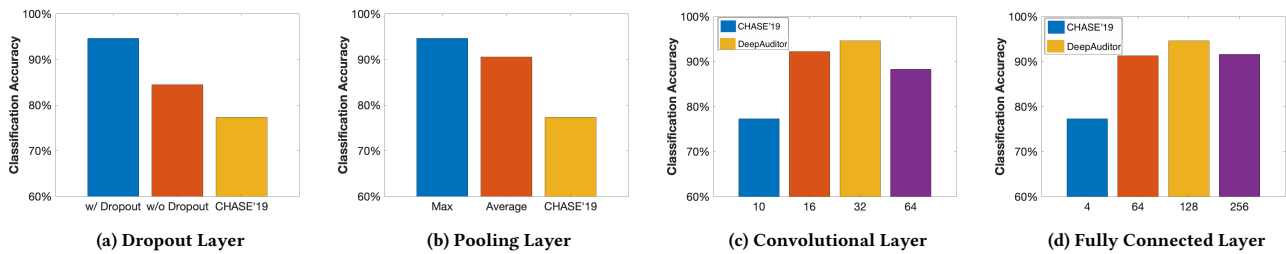


Figure 9: Impact of CNN Design Factors

5.1.2 Impact of Design Factors. In this vein, we compared our design factors with the previous model to show that our model outperforms it. In particular, we used the leave-router-out test to illustrate the performance difference. In this test, the training dataset includes security camera and voice assistant data, while the test dataset has the unseen router power-trace data. Figure 9 demonstrates the design factors of our 1-D CNN. First, Figure 9a shows that the dropout layer enhanced the validation accuracy from 84% to 94%. The dropout layer is helpful because it reduces the number of trainable features to prevent overfitting during training. The CHASE'19 model did not use a dropout layer, so the test and validation accuracy had a gap of 10% and the validation accuracy was below 80%. In Figure 9b, we observe that the max-pooling method enables to achieve a better classification accuracy than the average-pooling method. This result suggests that the max-pooling method can capture dramatic power-trace spikes well, while the average pooling method smoothes out those power spikes that may include important information. In Figure 9c, we achieved a better accuracy by increasing the number of kernels from 10 to 32. Note that using 64 kernels decreased the prediction performance because the unnecessary learning ability (more features) overfitted to the dataset. Lastly, Figure 9d illustrates that we determined the best number of neurons in the fully connected layer is 128. As illustrated, increasing the number of neurons to 256 hurt the validation accuracy, which hints at a possible overfitting problem.

Overall, we designed a powerful CNN classifier for botnet intrusion detection by optimizing design choices including hyperparameters. As a result, our CNN classifier outperforms the baseline model, which is too simple and overfitted to the dataset in leave-one-out tests.

5.1.3 Offline Validation of the Proposed Model. For practical deployments, neural network models need to be accurate against unseen patterns. Therefore, we further conducted offline experiments to determine the robustness of our model. Table 2 compares the DeepAuditor and CHASE'19 models.

Leave-one-device-out tests are meaningful because they show the robustness of the classifier in the practical deployment environment. For example, it is not possible to train power-trace data from all IoT devices because there are numerous IoT devices in the market. Thus, our classifier needs to be robust against unseen data

from new and future IoT devices. As shown in the previous subsection, in the leave-router-out tests, we trained power data from a security camera and a voice assistant. Then, the router dataset was used to test the classification accuracy. Our DeepAuditor CNN classifier achieved an accuracy of 94.6%, while the CHASE'19 model predicted power traces with an accuracy of 77.3%. In the leave-voice-assistant-out test, we also achieved a nearly 10% improvement in accuracy.

Leave-one-botnet-out tests are also crucial since they demonstrate the CNN's ability to detect new botnet attacks. As discussed in Section 2, Mirai and its variants utilized brute-force attacks, which are still adversaries' preferred options for intrusion into IoT devices. Although these attacks use slightly different procedures, they all follow the model of brute-force attacks followed by post-processing jobs. Thus, it is feasible to identify malicious intrusions of different botnets through power traces. We achieved over 95% prediction accuracy, which outperforms the baseline classifier's accuracy of below 80%. The above results clearly show that our model is generalized to different types of devices or botnets.

5.2 Privacy-preserved Inference Protocol

Based on our CNN classifier design, we propose a privacy-preserved inference protocol. This protocol was designed between the Data Inferencer and the Computing Cloud to offload computations without data leakage. We will theoretically validate the data protection of the protocol in Section 7.3.

5.2.1 Motivation and Design Principle. To offload computation resources, we separate multiplication and summation operations of the convolutional and fully connected layers into the two distributed components. In particular, Data Inferencer only computes cheap plaintext summations, while Computing Cloud carries out the multiplication operations in the ciphertext. Simultaneously, to make it more efficient, we encode our input data in both Data Inferencer and Computing Cloud following the convolution kernel's order rather than keeping the data as original format. By doing this, we also avoid requiring time-consuming permutation operations [26] and offline secret sharing strategies in both of the convolutional and fully connected layers. This is an improvement on other works [30, 33, 53]. After that, we feed the output directly into the next layer.

To make it possible, we utilize Packed Homomorphic Encryption (PHE) to allow Data Inferencer to encrypt the power trace data before uploading it to Computing Cloud. Then, the Computing Cloud runs the CNN computations on the encoded ciphertext. Thus, the Data Inferencer encodes multiple plaintext data elements into one

Table 2: Comparison of Leave-one-out tests

System	Metric	Leave-router-out	Leave-voice-assistant-out	Leave-Masuta-out
CHASE'19 [25]	Accuracy	77.3%	80.7%	79.6%
	Accuracy	94.6%	89.1%	95.5%
	Recall	98.2%	94.7%	98.4%
DeepAuditor	Precision	94.7%	87.4%	95.4%
	F1-Measure	96.4%	90.9%	96.9%

ciphertext and high-efficiently carries out element-wise homomorphic computations in a Single Instruction Multiple Data (SIMD) manner [8]. This tool is particularly useful for our system as our input instance includes thousands of sampling data due to the high sampling rate. Overall, our design uses the CKKS-based PHE that works on element-wise float point data addition and multiplication in ciphertext [12, 32].

Protocol 1 Privacy-preserving CNN Inference

Inputs: Received power trace X

Outputs: Prediction Status l on the given trace X

The Protocol:

1. Data Inferencer receives raw data X from Power Auditor, encrypts it as $[X']_C$, and sends $[X']_C$ to Computing Cloud.
2. Computing Cloud computes $[U]_C = K_1 W'_1 [X']_C + K_1 B'_1 + N_1$ where K_1 is a non-zero positive random vector, W'_1 is the encoded weight, B'_1 is the encoded bias, N_1 is a pseudo-random zero-sum vector. Finally, Computing Cloud sends $[U]_C$ to Data Inferencer.
3. Data Inferencer decrypts $[U]_C$ and computes $Z_j = \sum_{i=0}^3 U_{i+j}$, where j is the convolution block index. Then, it feeds the result to the ReLU activation function and max pooling layer. Finally, it gets the result Y .
4. Data Inferencer encodes and encrypts Y as $[Y']_C$, and then sends it to Computing Cloud.
5. Computing Cloud removes the random number K_1 and computes the $[V']_C = K_2 W'_2 [Y']_C + K_2 B'_2 + N_2$, where the K_2 is a non-zero positive random number, N_2 is a pseudo-random zero-sum vector, and the subscript 2 of W and B indicates the corresponding variables at the fully connected layer in our CNN model. Then, Computing Cloud sends $[V']_C$ to Data Inferencer.
6. Data Inferencer decrypts $[V']_C$ and performs the summation similarly to Step 3 on each hidden unit block. Finally, it feeds the fully connected layer output into the softmax layer to get the final prediction status l for a given power trace. Depending upon the policy, Data Inferencer takes further actions if l represents a malicious intrusion.

5.2.2 Detailed Procedures. The detailed privacy-preserved inference protocol is described in Protocol 1 and Figure 10. In order to better explain the key idea, we take an identical CNN model with a smaller size input instead of the original input size. Recall that the input size of our CNN model is (1×2550) , and the kernel size is (1×128) with the stride size 64 (See Section 5.1). Instead, we use a CNN model as an example whose kernel size is (1×4) with the stride size 2.

Let X denote the received raw data by Data Inferencer from Power Auditor. The Data Inferencer first utilizes a PHE package that uses one packed vector to store multiple encrypted plaintext data. In Step 1, to implement the convolution function over the ciphertext, Data Inferencer encodes the data X to X' , as illustrated in Figure 10. Correspondingly, Computing Cloud encodes the weight W_1 and bias B_1 into packed vectors W' and B' in Step 2. With such

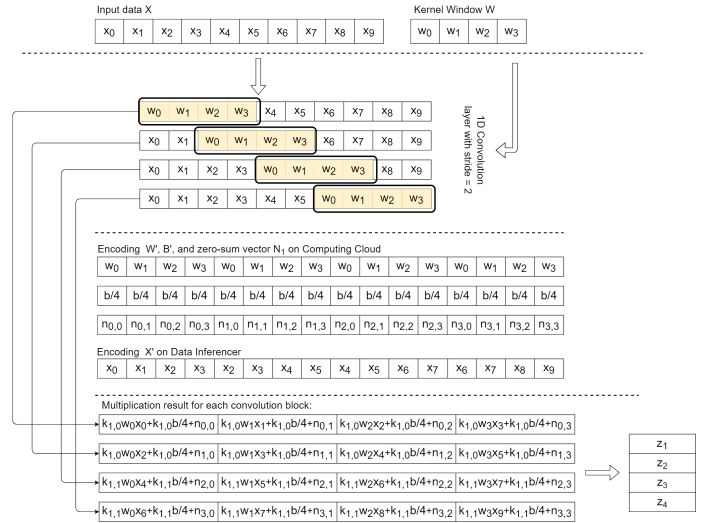


Figure 10: Data operations on Data Inferencer and Computing Cloud in Steps 2 and 3

encoding, the convolution between X and W can be implemented as the element-wise multiplication between X' and W' , plus B' .

Steps 2 and 3 show how we securely implement the convolutional layer among ciphertext. After Computing Cloud receives the ciphertext $[X']_C$ from Data Inferencer, Computing Cloud uses Eq. (1) to compute the homomorphic multiplication result.

$$[U]_C = K_1 \times W'_1 \times [X']_C + K_1 \times B'_1 + N_1, \quad (1)$$

The purpose of using random numbers N_1 and K_1 in Eq. (1) is to prevent Data Inferencer from inferring the model parameter W'_1 from its received message $[U]_C$. Computing Cloud first generates a zero-sum vector $N_1 \in \mathbb{Z}$, which is a vector of pseudo-random numbers such that $N_1 = \sum_{j=0}^3 n_{1,j} = 0$ ($0 \leq j \leq 3$), to mask each multiplication result, as illustrated in Fig. 10. Then, Computing Cloud multiplies a non-zero positive random number vector $K_1 = [k_{1,0}, k_{1,1}]$ to mask all multiplication results. With both masks N_1 and K_1 , Data Inferencer is unable to learn the parameter W'_1 and B'_1 based on $[U]_C$ and X' . Note that N_1 is different across different kernels in convolution. Simultaneously, recall that the pool size of max pooling size is 2, K_1 is different for every two kernels to make sure the correctness for the output of max pooling layer.

Steps 5 and 6 are similar to Steps 2 and 3 but implement the fully connected layer. However, Computing Cloud only requires to choose a non-zero positive random number K_2 in Step 5 to mask the ciphertext multiplication result. At the end of Step 6, Data Inferencer directly feeds the weighted sum result of the fully connected layer into the softmax layer to infer the Power Auditor status l . If the l value represents a malicious intrusion on the IoT Device, Data Inferencer can take further actions, such as sending a notification to the administrator or shutting down the IoT device.

6 ONLINE SYSTEM IMPLEMENTATION

For performance evaluation, we have implemented a prototype system in a distributed setting. Section 6.1 describes the prototype implementation and Section 6.2 shows the new dataset we collected.

Table 3: Online System Testbed

	Power Auditor	IoT Device	Mirai Bot	Data Inferencer	Computing Cloud
Network Deployment	Local Network				Cloud
Hardware Platform	Raspberry Pi 3	Raspberry Pi 3	Raspberry Pi 3	Apple Mac Mini Desktop	Linux Server
CPU	1.4GHz, Quad-core	1.4GHz, Quad-core	1.4GHz, Quad-core	2.5GHz, Dual-core	2.6GHz, 2*16 cores
Memory	1GB	1GB	1GB	8GB	250GB
Sensor used	Current Sensor (INA219A)	- Camera Sensor - External Microphone	N/A	N/A	N/A
Operating System	Raspbian 9	Raspbian 10	Raspbian 9	macOS 10.15	Ubuntu 18.04
Software Module	- Power Monitoring - Networking - Power Controlling	- Video Streaming - Motion Detection - Voice Assistant	Mirai (Scanner/Loader)	Privacy-preserved Online Classifier	Privacy-preserved Online Classifier

6.1 DeepAuditor Implementation

We have built a prototype DeepAuditor for proof of concept. Table 3 describes our testbed environment. First, Raspberry Pi 3 devices serve as both Power Auditors and IoT devices for prototyping purposes. We also used a desktop in the same local network for Data Inferencer. Computing Cloud was deployed in a department cloud server outside the local network.

In Power Auditors, we implemented the proposed modules in Python for real-time data collection. For prototyping IoT devices, we installed an open-source camera software called MotionEye [9] on the connected IoT devices. This software includes a motion detection feature as well as a video streaming feature. Thus, we consider the Raspberry Pi 3 device running MotionEye as an IoT device. Likewise, we also installed the Google AIY project [20] on another Raspberry Pi 3 device and conducted voice commands on it. The Power Auditor is connected to an AC adapter. The IoT device is then supplied with power through the Power Auditor, as shown in Figure 3.

Next, we implemented the sliding window scheme between the Power Auditor and the Data Inferencer. Since the Power Auditor sends segmented power traces of 1.5-second data every 0.5 seconds, the Data Inferencer is required to predict the power trace data every 0.5 seconds.

As mentioned in Section 5.2, our privacy-preserved protocol utilized CKKS-based PHE [12], which involves three parameters¹: 1) polynomial modulus degree N , 2) ciphertext scale s , and 3) coefficient modulus. In our protocol, we set up the CKKS-related encryption parameters as follows: 1) Parameters for PHE scheme are selected for a 128-bit security level, 2) the selection of polynomial modulus degree N is a smaller degree that allows to encode $N/2$ elements into one ciphertext. In our case, we selected $N = 32768$ for the convolution layer and $N = 4096$ for the fully connected layer, 3) a ciphertext scale $s = 2^{40}$ is enough to store all the intermediate results in the convolutional layer, and a ciphertext scale $s = 2^{20}$ is enough to store all the intermediate results in the fully connected layer, and 4) In seal [32], the modulus switching chain is set up as coefficient modulus for ciphertext against exponential noise growth during ciphertext operations. In our case, we selected the coefficient modulus (60, 40, 40, 60) for the convolutional layer and (30, 20, 20, 30) for the fully connected layer.

With the above parameters, we fulfilled the CNN classifier in the Data Inferencer and the Computing Cloud. According to Protocol 1, the classification procedures comprise six steps; steps 1, 3, 4, and 6 are implemented in the Data Inferencer, whereas steps 2 and 5 are

¹The reader is referred to [12, 32] for more details.

Table 4: The Collected Dataset of Power Traces

Class	Description	Number of Instances	Total number of Instances
Idle	When IoT Service is not running	4693	4693
IoT Device Service	Security Camera [9]	5976	8176
	Voice Assistant [20]	2200	
Reboot	When system is rebooting	2200	2200
Botnet (Mirai) [22]	Malicious behavior while system is Idle	2000	3000
	Malicious behavior while system is running IoT service	1000	

implemented in the Computing cloud. Step 6 makes a final decision for power-trace prediction.

To satisfy the real-time prediction requirement, we adopted pipeline processing [45] in the six inference steps; each step is being executed in parallel in the Data Inferencer and the Computing Cloud. In other words, the Data Inferencer and the Computing Cloud do not have to wait until all the steps are completed. This pipelining increases the throughput of the instructions, which eventually leads to our real-time inference for multiple IoT devices simultaneously. We will discuss the performance in Section 7.4.

6.2 Dataset Collection for Online Test

As discussed earlier, our classifier predicts which of the four classes a power instance belongs to. Even though the previous section already demonstrated the robustness of our classifier design on the public dataset, we newly collected power traces from two different types of IoT devices to demonstrate the real-time performance. Thus, we created a new dataset in our testbed environment, as shown in Table 4.

Table 4 summarizes the collected dataset. In our environment, we generated a specific scenario and collected over 2,000 power instances for each class. For example, the data we collected for the Idle class comprises 4,693 instances of 1.5-second power traces when the IoT service was not running. We also collected power traces when the IoT service was running or the IoT device was rebooting. For example, we used the open-source MotionEye [9] for security cameras and the Google AIY project [20] for voice assistant prototypes. For the Botnet class, we downloaded an open-source code of Mirai from Github [22] and built it on an IoT bot testbed. To generate Mirai instances in our local network, we modified the source code to attack only our IoT devices. Then, we collected 3,000 instances of malicious attacks when the IoT service is running or the system is idle.

7 ONLINE SYSTEM EVALUATION

In Section 7.1, we validate the performance of the Power Auditor device. Section 7.2 then demonstrates the system-level online classification performance. In Section 7.3, we theoretically analyze the data protection of the privacy-preserved inference protocol. Finally, Section 7.4 illustrates the scalability evaluation of the system.

7.1 Power Auditor Performance

In Table 5, we compare our power auditing device with the off-the-shelf device Monsoon Power Monitor [35] that was used in

Table 5: Power Auditing Devices Comparison

	Our Power Auditor	Monsoon Power Monitor [35]
Sampling Rate	Up to 1.7kHz	Up to 5kHz
Measurement Range	Up to 5.5V Up to 2.3A	Up to 13.5V Up to 6A
Dimension	3" x 2" x 1"	8" x 6" x 3"
Weight	0.3lb	4lb
Software	Text-based Python Software	Window GUI Software
Online Measurement	Available	Not Applicable
Price	\$25	\$929

the offline study [25]. Our Power Auditor supports a sampling rate of up to 1.7kHz and a voltage of up to 5.5V. These ranges are less than the Monsoon device and thus may be limited in measuring the power of large appliances with built-in computation units, e.g., smart fridges or smart microwaves. However, it is still enough to audit most IoT devices' power consumption. On the contrary, the Power Auditor is much smaller and lighter than the Monsoon device, which makes our device convenient for ubiquitous power measurement. More importantly, the proposed device is capable of measuring power-trace data in real-time for online inference. The Power Auditor also supplies power to connected devices while measuring power consumption.

Table 6 illustrates the performance metrics for a single Power Auditor working with our servers. In short, the resources of Raspberry Pi 3 are more than capable of supporting the Power Auditor device. For example, the maximum CPU load in the Power Auditor is up to 76%, which used roughly 20% of the Raspberry Pi's quad-core computing power. Memory (RAM) usage during online auditing is only 10MBytes out of Raspberry Pi's 1GB (1%). Based on the proposed sliding window design, the required network bandwidth between the Power Auditor and the Data Inferencer is only 120Kbps. This bandwidth is extremely small and therefore can be covered by Bluetooth or even the Zigbee protocol. Moreover, the power consumption of the Power Auditor is approximately 2W during the real-time inference. That is about an extra \$2 in cost per year for single-device monitoring, which is similar to the power consumption of existing smart plug devices [3]. Note that these results were the same regardless of IoT device type. Due to space constraints, we did not include those results. Overall, our results confirm that the Power Auditor is lightweight enough for an online auditing device, and we plan to build a prototype upon real-world demonstration.

7.2 Online CNN Classifier Performance

We measured online classification results in a laboratory setting. We provide the classification accuracy and other metrics associated with this test to validate the classifier performance. We generated power instances of each class in real-time and conducted online inference to obtain metrics. For example, the IoT device was in idle

Table 6: Online System Performance per IoT Device

	CPU Load (Max ¹)	Memory Usage	Network Bandwidth	Processing Delay	Power Consumption
Power Auditor	76% (400%)	10MB	120Kbps	25ms	2W (400mA)
Data Inferencer	40% (200%)	50MB	11.375Kbps	160ms	—
Computing Cloud	35% (3200%)	30MB	—	360ms	—

¹ Maximum CPU Load depends on the number of CPU cores, e.g., Dual-core has a maximum 200% CPU load.

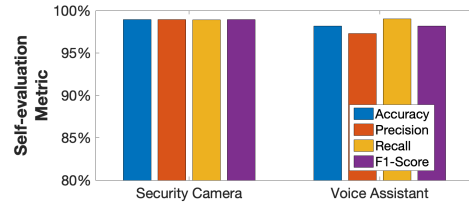


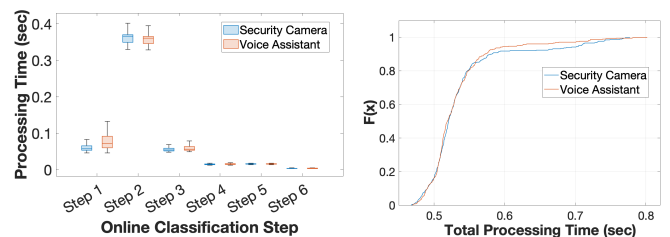
Figure 11: Online Classification Results

status for the Idle class, while the IoT device streamed video or conducted voice commands for the IoT Service class. For the Mirai intrusion class, a bot device sneaked into the IoT device, and we measured the inference results on the intrusion events.

Table 6 shows the performance of each distributed component as the classifier is deployed in separate servers. These metrics were obtained when the DeepAuditor processed continuous data from a single Power Auditor monitoring an IoT device. The total delay per power instance is approximately 520 milliseconds: 160ms in the Data Inferencer and 360ms in the Computing Cloud. In our testbed, the transmission delay between the Data Inferencer and the Computing Cloud is approximately 20ms, while the local network delay is less than 1ms. The network bandwidth between the Data Inferencer and the Computing Cloud is 11.375Kbps, and the memory usage is 50MBytes in the Data Inferencer and 30MBytes in the Computing Cloud. Overall, these numbers are not overwhelming for online inference since we utilized cloud resources.

Figure 11 also shows the online classification results for each IoT device. The results demonstrate the exceptional classification ability of the CNN classifier. For example, we achieved an overall accuracy of 98.95%. The Precision and Recall metric values for both tests are also above 98%. Thus, the distributed classifier is able to distinguish different patterns of device behavior in real-time, as trained.

Furthermore, Figure 12a illustrates the processing delay of each step in the privacy-preserved inference protocol. For instance, it takes about 350ms to complete the Convolution procedure (Step 2) on Computing Cloud regardless of device type, which is the majority of the entire online classification. Other steps consume relatively fewer computing resources. In addition, Figure 12b demonstrates an empirical CDF function of the online classification response time. For both devices, we observe that over 80% of the inferences were done in 550ms or less. Note that our system is required to classify real-time instances per Power Auditor every 500ms. To maximize the process throughput, we applied pipeline processing to the inference protocol. Thus, the entire performance mainly rests



(a) Processing Time per Protocol Step (b) Empirical CDF of Response Time

Figure 12: Online Classifier Processing Time

on step 2, which is the most time-consuming job in our classifier. Even though most instances are completed about 550ms after the Data Inferencer receives the power instance, the Computing Cloud is able to classify two instances per second securely because step 2 takes at most 400ms, as presented in Figure 12a. Overall, the DeepAuditor as an entire system is able to reliably predict input instances every 500ms in real-time.

7.3 Theoretical Analysis of the Privacy-preserved Inference Protocol

In this subsection, we demonstrate that our distributed classifier design is secure in that 1) Computing Cloud cannot obtain the client's power-trace data, and 2) Data Inferencer cannot obtain the model parameters W and B of the CNN model in Computing Cloud. Hence, there is no information leakage between Computing Cloud and Data Inferencer.

We used a security analysis method called *ideal/real world* paradigm [38]. Let P1 denote Data Inferencer with input x and P2 Computing Cloud with input y . Let $f = (f_1, f_2)$ be a set of functionality and π be a protocol which is implemented in our online system for computing f . The party P_i wishes to obtain the protocol output $f_i(x, y)$ ($i \in \{1, 2\}$). The view of P_i during an execution of π on the set of input $\bar{x} = \{x, y\}$ is denoted as $view_i^\pi(\bar{x})$, and it equals to $(w, r_i; m_1^i, \dots, m_t^i)$ where $w \in x, y$ is the input of P_i for $i \in \{1, 2\}$. r_i is equals to the set of random numbers inside P_i , and m_i^j represents the j -the message received by P_i .

The adversary can compromise any one party in our model, but not the majority and all of them are non-colluded. There exists a probabilistic polynomial-time simulator that has same functionality as π . The security of π is defined as follows:

Definition 7.1. A Protocol π can securely execute f in the presence of semi-honest adversaries, if for any exists a probabilistic polynomial-time simulator S_i , such that for the corrupted parties P_i , it has:

$$\begin{aligned} \{(S_1(1^n, x, f_1(\bar{x})), f(\bar{x}))\} &\equiv view_i^\pi(\bar{x}) \\ \{(S_2(1^n, y, f_2(\bar{x})), f(\bar{x}))\} &\equiv view_i^\pi(\bar{x}) \end{aligned} \quad (2)$$

where \equiv represents computationally indistinguishable, and $i \in \{1, 2\}$.

The protocol π is secure against semi-honest adversaries if the views of the real-world execution are computationally indistinguishable from the view of the simulator in ideal world. In the following section, we use the method above to prove our online system is secure against semi-honest adversaries.

7.3.1 Data Security against Compromised Data Inferencer. We assume that Data Inferencer is compromised by Adversary A. We use the simulator sim to behave as Adversary A which interacts with f . The sim , f , and Data Inferencer conduct the following steps:

1) sim first gets $[X']_C$ from Data Inferencer. Then, it sends $[\bar{X}']_C$ to f . f returns $[U]_C$ to sim .

2) Starting Data Inferencer, sim generates random number \bar{W} , \bar{B} , \bar{K} and \bar{N} , and computes $[U']_C = \bar{K}\bar{W}[X']_C + \bar{K}\bar{B} + \bar{N}$ by Eq. (1). Then, sim sends $[U']_C$ to Data Inferencer.

3) sim decrypts and outputs U and U'

The above protocol is secure against adversary Data Inferencer based on the randomness of K' and N' which makes the view of U' that is generated by sim computationally indistinguishable from the view of U that the real output from f . This conclusion can also be applied to the view of the Fc layer in our system.

7.3.2 Data Security against Compromised Computing Cloud. Similarly, we assume that Computing Cloud is compromised by Adversary A which interacts with sim . The sim , f , and Computing Cloud conduct the following steps:

1) sim first gets $W_1, B_1, K_1, N_1, W_2, B_2, K_2$, and N_2 from Computing Cloud. Then, it sends $W_1, B_1, K_1, N_1, W_2, B_2, K_2$, and N_2 to f . f returns None to sim .

2) Starting Computing Cloud, sim generates and encrypts a group of random numbers \bar{X} as $[\bar{X}]_C$. Then, sim sends it to Computing Cloud.

3) sim receives $[U']_C$ from Computing Cloud, sim generates and encrypts a group of random numbers \bar{Y} as $[\bar{Y}]_C$. Then, sim sends $[\bar{Y}]_C$ to Computing Cloud.

4) sim outputs $([\bar{X}]_C, [\bar{Y}]_C)$.

Our system is secure against Computing Cloud because the view of the Data Inferencer's input data $[X]_C$ and the intermediate output $[Y']_C$ is computationally indistinguishable from the $[\bar{X}]_C$ and $[\bar{Y}]_C$ that is generated by sim based on the fact that the PHE algorithm is semantically secure [8].

7.3.3 Computation Complexity Analysis. We analyzed the overall computation complexity for the convolutional layer and fully connected (FC) layer of DeepAuditor in Table 7. Let denote that r is kernel size for the convolutional layer and C is the number of output channels for the convolutional layer. We compared our work with a naive method of Gazelle [26], which is a more state-of-the-art protocol with speed-up than some classic privacy-preserving inference protocol like [19]. Based on our benchmark on Protocol 1, Computing Cloud needs C times ciphertext multiplication and addition to compute the intermediate result $[U]_C$. Meanwhile, Data Inferencer only conducts a cheap plaintext summation $[U]_C$ to complete the convolution output Z_j . The total computation cost for our protocol design is much less than the Gazelle, which requires r times ciphertext permutation, rC times ciphertext multiplication, and rC times ciphertext addition.

To improve the computation efficiency in the FC layer. We extended our data encode operation presented in Figure 10, which enables all max-pooling outputs to be packed into one ciphertext. Let n_i denote the number of data in each output channel of max-pooling, n_o denote the output data for the FC layer, n denote the number of slots for one ciphertext, and k denote the kernel size for the FC layer. Each ciphertext can hold $\frac{Cn_in_o}{n}$ data. Compared with Gazelle's input packing method [26], the Computing Cloud in our protocol needs k times less than ciphertext multiplication and

Table 7: Comparison of Computation Complexity

Methodology	Permutation	Multiplication	Addition
Gazelle-Convolution	$O(r)$	$O(rC)$	$O(rC)$
DeepAuditor-Convolution	0	$O(C)$	$O(C)$
Gazelle-FC	$O(k)$	$O(k \frac{Cn_in_o}{n})$	$O(k \frac{Cn_in_o}{n})$
DeepAuditor-FC	0	$O(\frac{Cn_in_o}{n})$	$O(\frac{Cn_in_o}{n})$

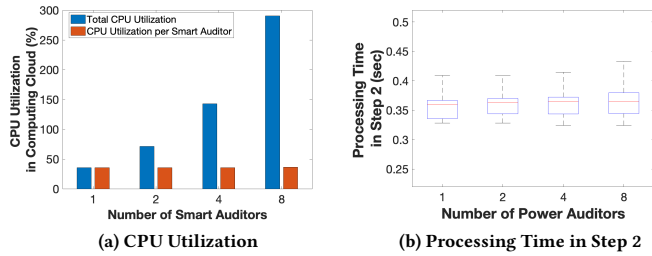


Figure 13: Scalability Performance Evaluation

addition. The FC layer of our protocol also does not require ciphertext permutation, while Gazelle still requires k times ciphertext permutation.

7.4 Scalability Evaluation

As shown in Figure 2, a single Computing Cloud supports multiple Power Auditors. We deployed the online CNN classifier in a distributed environment to offload computation resources and handle multiple IoT devices simultaneously. In our testbed, we evaluated how many IoT devices the cloud servers could support for intrusion detection.

To test scalability, we set up a distributed environment with eight Power Auditors, two Data Inferencers, and one Computing Cloud. Four Power Auditors are connected to each Data Inferencer, both of which are connected to the same Computing Cloud. The eight Power Auditors collect power traces from their respective IoT devices in real-time. Note that system scalability relies on efficient use of available resources, such as throughput and CPU utilization [24], regardless of device or botnet type. Figure 13 depicts performance results of four tests of 1, 2, 4, and 8 Power Auditors. In Figure 13a, as the number of Power Auditors increases, the CPU utilization of the Computing Cloud increases linearly. The Computing Cloud’s CPU utilization per Power Auditor is approximately 35%, which is consistent with the result in Table 6. Moreover, when we tested with eight Power Auditors, the total CPU utilization was less than 300% out of 3,200% (32 Cores).

Furthermore, Figure 13b demonstrates that the DeepAuditor classifies multiple IoT devices’ data in real-time. As discussed in Section 7.2, the inference time mostly relies on the processing time in step 2 of our protocol. In Figure 13b, as we increased the number of Power Auditors, the average processing time does not change substantially. As long as the processing time in step 2 is less than 0.5 seconds, our DeepAuditor system can guarantee real-time inference for online detection.

Overall, these results demonstrate that the Computing Cloud supports inference computations for multiple IoT devices to the extent that CPU cores are available. For example, step 2 takes approximately 350ms on average, and the CPU utilization of Computing Cloud is about 35% for handling a single Power Auditor. Thus, a single core of Computing Cloud can handle up to three Power Auditors per second. Since the Computing Cloud in our environment has 32 CPU cores, this server can support almost a hundred Power Auditors. We further plan to enhance the convolution processing time in the future because the current performance is bounded by the processing time in step 2.

Table 8: Our classifier was validated on different datasets

	Experimental Environment	Power Monitoring Device	Year Collected	Number of Device Type	Number of Instances
CHASE’19 [1]	Lab	Monsoon Power Monitor	2019	3	8000
DeepAuditor	At-home	Our Proposed Device	2021	2	15869

8 DISCUSSION

In this section, we present our thoughts with regard to limitations and future work.

8.1 Power Auditor Prototype

Since we utilized a Raspberry Pi device for prototyping the Power Auditor, our current version is still bulky and costly for ubiquitous power measurement. Several researchers have already developed a small form-factor power meter for communication purposes [23] [16]. However, because it was not designed for IoT devices, the power ranges are slightly higher than those of many IoT devices, whereas our current Power Auditor is able to monitor low-powered IoT devices. Thus, we plan to make our Power Auditor an AC-plug meter device that can be attached to off-the-shelf IoT devices. This task requires PCB board manufacturing and is shown to be feasible by other works. If so, the Power Auditor will be even smaller and cheaper than the current prototype. Moreover, since the smart plug industry has grown dramatically, we believe that in the future, our model can be integrated into generic smart plug devices for online intrusion detection.

8.2 Real-world Deployment

DeepAuditor is the first distributed online system to assess the power consumption of IoT botnet intrusion following the emergence of Mirai and similar IoT botnets. Fundamentally, our distributed system components process the power consumption of the connected IoT devices in real-time for intrusion detection, which has never been tackled before. Thus, we focused on demonstrating the validity of our system design since other approaches may not be directly comparable with the DeepAuditor. Nevertheless, our proposed CNN classifier was tested with two different datasets. Table 8 shows the environments of the two datasets, including location, time, and power monitor devices. As demonstrated in the previous sections, our classifier performed well in both environments.

However, our system still needs to be tested in a real-world deployment. Based on the proven concept, we plan to further expand our current system to a wild setting in which commercial IoT devices will be used for validation purposes. We will tackle this task in the future.

9 RELATED WORK

Section 9.1 presents side-channel studies on IoT security. In Section 9.2, we summarize the related work with regard to IoT security via power auditing. Then, Section 9.3 introduces the existing work concerning about the preservation of data privacy.

9.1 IoT Security via Side-channel Information

The literature has scrutinized IoT security against botnets for decades [7, 49]. This area includes network-based solutions as well as power-auditing-based detection methods. While network-based solutions have struggled to address the endpoint security on IoT devices [10], several works have utilized side-channel information, such as electromagnetic (EM) data. Especially for resource-constrained devices, using side-channel information is efficient because it utilizes existing resources instead of requiring many modifications [39].

Nazari et al. [37] proposed an EM-based spike detection method to identify injected codes in program execution. The EM spectrum was monitored in order to detect malware, and the results showed promise in using side-channel information. Sehatbakhsh et al. [42] also utilized EM-based monitoring to identify anomalous behavior during execution on medical devices. While EM-based side-channel monitoring also showed potential in program execution, system-level monitoring is preferred in the identification of malicious IoT botnets. This is because IoT botnets often enter and compromise entire target devices [25].

9.2 IoT Security via Power Auditing

Power side-channel information has also been utilized to infer malicious behavior on end devices. For example, some pioneering works used power side-channel data to detect malign behavior on mobile devices in the early 2010s [27][51]. Recently, several researchers [25, 29, 36] have worked on IoT devices to characterize malicious behavior as IoT botnets have been popular. Myridakis et al. [36] implemented a power monitoring circuit for botnet prevention for IoT devices. However, this system mainly focused on detecting massive DoS attacks on IoT devices with a spike detection method instead of intrusion detection. Li et al. [29] addressed energy auditing for physical and cyber attacks. For cyber attacks, it utilized dual-CNNs to infer massive DoS attacks such as network flood using energy meters [21]. Clark et al. [14] aimed to identify malware behavior on medical devices via power auditing. Similarly, they conducted offline classification experiments on a dataset collected from a single device. Jung et al. [25] pioneered IoT botnet intrusion detection via power modeling. While their CNN classifier showed promise in detecting subtle differences in power traces, the study was conducted offline with a bulky and expensive power monitor. Thus, it is not practical for ubiquitous botnet detection on IoT devices. In summary, Table 9 summarizes the related works on IoT security via power auditing. Overall, there is still a gap between power auditing techniques and practical IoT intrusion detection. For efficient IoT botnet detection, a scalable real-time solution via power auditing is needed. We are the first to realize a distributed online classification system for botnet intrusion detection on multiple IoT devices via ubiquitous power auditing.

9.3 Preservation of Data Privacy

Preservation of data privacy has been widely studied in the literature. There are three major approaches. The first approach is differential privacy [17], which injects noise into query results, such as perturbing stochastic gradient descent (SGD) [4]. However, the additive noise may degrade model accuracy. The second approach designs privacy-preserved protocols based on secure multi-party

Table 9: Comparison of IoT Security via Power Auditing

	Target Attacks	Testbed Environment	Learning Method	Concurrent Capacity	Auditing Device	Classification Accuracy
Jung [25]	Botnet Intrusion	Offline Modeling	1-D CNN	Single Device	Monsoon [35]	96.5%
Li [29]	DoS Attacks	Online Classification	Dual 1-D CNNs	Single Device	IoT Hardware [21]	MSE 0.032
Myridakis [36]	DoS Attacks	Online Classification	Spike Detection	Single Device	IoT Hardware	100%
Clark [14]	Malware Attacks	Offline Classification	kNN, RF, Perceptron	Single Device	AC Outlet [13]	94%
DeepAuditor	Botnet Intrusion	Online Classification	Distributed 1-D CNN	90+ Devices	IoT Hardware	98.9%

computations. They usually distribute secrets among a group of parties to achieve security computations at the expense of high computational overhead and strong security assumptions [11][52][34]. Thus, they are rarely adopted in general scenarios.

A new solution for privacy preservation was introduced by using the fully homomorphic encryption [18]. It allows users to encrypt data and offload the computation to a cloud. The cloud computes encrypted data and sends back encrypted results [19][30][26]. However, the nonlinear activation computation cannot be supported by the homomorphic encryption, and the approximation often has to be used. Compared with existing work, our solution is novel in that our proposed scheme capitalizes on the proposed CNN model structure to adopt a smart design to address this problem.

10 CONCLUSION

In this paper, we proposed a distributed online intrusion detection system for IoT devices via power auditing. We first developed a portable power-auditing device to measure power side-channel information of IoT devices in real-time. The one-dimensional CNN classifier was then designed and deployed in a distributed setting. The online CNN classifier predicted IoT devices' behavior with up to 98.9% accuracy, which outperforms the baseline classifier, especially in leave-one-out tests. In addition to the system components, we also designed distributed protocols to avoid data leakage and reduce networking redundancy. Finally, we evaluated the scalability of the system in a laboratory setting. Altogether, the DeepAuditor system is the first online intrusion detection system that classifies multiple IoT devices' behavior via power traces. This kind of cloud system that uses power auditing of multiple IoT devices has not been addressed before in the literature, so such system can be used on IoT devices for other purposes.

In the future, we plan to enhance the performance of the inference protocol. Currently, the convolutional layer consumes the majority of the entire processing time. If we reduce that procedure, our system will be more reliable and scalable. In addition to the pre-trained classifier, we further plan to apply unsupervised learning so that users can use their dataset without labeling. This can expedite system deployment in a practical setting.

ACKNOWLEDGEMENTS

This research is partially supported by COVA CCI Cybersecurity Research and Innovation Funding, COVA CCI Cybersecurity Innovation Bridge Fund (Grant #HC-4Q21-005), COVA CCI Dissertation Fellowship, and NSF grant CNS-2120279. We would like to thank all the anonymous reviewers for their valuable comments.

REFERENCES

- [1] Iot-botnet-detection via power consumption modeling. <https://woosup.github.io/IoT-Botnet-Detection>.
- [2] Smart plug market - growth, trends, covid-19 impact, and forecasts (2021 - 2026). <https://www.mordorintelligence.com/industry-reports/smart-plug-market>.
- [3] What is a smart plug and how it eliminates energy waste. <https://www.atlanticeenergyco.com/post/smart-plug-energy-savings/>.
- [4] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, pages 308–318, New York, NY, USA, 2016. ACM.
- [5] Adafruit. *INA219 Current Sensor*, 2020. <https://learn.adafruit.com/adafruit-ina219-current-sensor-breakout>.
- [6] M. Antonakakis. Understanding the Mirai Botnet. *USENIX Security Symposium*, July 2017.
- [7] L. Aversano, M. L. Bernardi, M. Cimitile, and R. Pecori. A systematic review on Deep Learning approaches for IoT security. *Computer Science Review*, 40:100389, Jan. 9999.
- [8] Z. Brakerski, C. Gentry, and S. Halevi. Packed ciphertexts in lwe-based homomorphic encryption. In K. Kurosawa and G. Hanaoka, editors, *Public-Key Cryptography – PKC 2013*, pages 1–13, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [9] ccrisan. *MotionEye*, 2020. <https://github.com/ccrisan/motioneye>.
- [10] N. Chaabouni, M. Mosbah, A. Zemmar, C. Sauvignac, and P. Faruki. Network intrusion detection for iot security based on learning techniques. *IEEE Communications Surveys Tutorials*, 21(3):2671–2701, 2019.
- [11] T. Chen and S. Zhong. Privacy-preserving backpropagation neural network learning. *IEEE Transactions on Neural Networks*, 20(10):1554–1564, Oct 2009.
- [12] J. H. Cheon, A. Kim, M. Kim, and Y. Song. Homomorphic encryption for arithmetic of approximate numbers. In T. Takagi and T. Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017*, pages 409–437, Cham, 2017. Springer International Publishing.
- [13] S. S. Clark. *The security and privacy implications of energy-proportional computing*. University of Massachusetts Amherst, 2013.
- [14] S. S. Clark, B. Ransford, A. Rahmati, S. Guineau, J. Sorber, W. Xu, and K. Fu. {WattsUpDoc}: Power side channels to nonintrusively discover untargeted malware on embedded medical devices. In *2013 USENIX Workshop on Health Information Technologies (HealthTech 13)*, 2013.
- [15] CNet. These smart plugs are the secret to a seamless smart home, 2019.
- [16] S. DeBruin, B. Ghena, Y.-S. Kuo, and P. Dutta. Powerblade: A low-profile, true-power, plug-through energy meter. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 17–29, 2015.
- [17] C. Dwork, A. Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [18] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing, STOC '09*, pages 169–178, New York, NY, USA, 2009. ACM.
- [19] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International Conference on Machine Learning*, pages 201–210, 2016.
- [20] google. *Google AIY Projects*, 2021. <https://aiyprojects.withgoogle.com>.
- [21] A. Hindle, A. Wilson, K. Rasmussen, E. J. Barlow, J. C. Campbell, and S. Romansky. Greenminer: A hardware based mining software repositories software energy consumption framework. *dl.acm.org*, 2014.
- [22] jgamblin. *Mirai-Source-Code*, 2017. <https://github.com/jgamblin/Mirai-Source-Code>.
- [23] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity ac metering network. In *2009 International Conference on Information Processing in Sensor Networks*, pages 253–264. IEEE, 2009.
- [24] P. Jogalekar and M. Woodside. Evaluating the scalability of distributed systems. *IEEE Transactions on Parallel and Distributed Systems*, 11(6):589–603, 2000.
- [25] W. Jung, H. Zhao, M. Sun, and G. Zhou. IoT Botnet Detection via Power Consumption Modeling. In *ACM/IEEE CHASE*, 2019.
- [26] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan. {GAZELLE}: A low latency framework for secure neural network inference. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 1651–1669, 2018.
- [27] H. Kim, J. Smith, and K. G. Shin. Detecting energy-greedy anomalies and mobile malware variants. *MobiSys*, page 239, 2008.
- [28] M. Kuzin, Y. Shmelev, and V. Kuskov. New trends in the world of iot threats. <https://securelist.com/new-trends-in-the-world-of-iot-threats/87991/>, 2018.
- [29] F. Li, Y. Shi, A. Shinde, and J. Ye. Enhanced cyber-physical security in internet of things through energy auditing. *ieeexplore.ieee.org*, 2019.
- [30] J. Liu, M. Juuti, Y. Lu, and N. Asokan. Oblivious neural network predictions via minionn transformations. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 619–631, 2017.
- [31] A. Maiti and M. Jadhwal. Light Ears - Information Leakage via Smart Lights. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2019.
- [32] microsoft Research. *Microsoft seal (release 3.2)*, Feb. 2019. <https://github.com/Microsoft/SEAL>.
- [33] P. Mishra, R. Lehmkuhl, A. Srinivasan, W. Zheng, and R. A. Popa. Delphi: A cryptographic inference service for neural networks. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, pages 2505–2522, 2020.
- [34] P. Mohassel and Y. Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 19–38, May 2017.
- [35] Monsoon. Monsoon power monitor. <https://www.msoon.com/high-voltage-power-monitor>, 2017.
- [36] D. Myridakis, P. Myridakis, and A. Kakarountas. A Power Dissipation Monitoring Circuit for Intrusion Detection and Botnet Prevention on IoT Devices. *Computation*, 2021.
- [37] A. Nazari, N. Sehatbakhsh, M. Alam, A. Zajic, and M. Prvulovic. Eddie: Em-based detection of deviations in program execution. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pages 333–346, 2017.
- [38] G. Oded. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, USA, 1st edition, 2009.
- [39] O. Or-Meir, N. Nissim, Y. Elovici, and L. Rokach. Dynamic malware analysis in the modern era—a state of the art survey. *ACM Computing Surveys (CSUR)*, 52(5):1–48, 2019.
- [40] N. Panwar, S. Sharma, S. Mehrotra, L. Krzywiecki, and N. Venkatasubramanian. Smart Home Survey on Security and Privacy. *arXiv.org*, 2019.
- [41] Radware. A game of cat and mouse: Dynamic ip address and cyber attacks, Feb. 2016. <https://security.radware.com/ddos-threats-attacks/ddos-attack-types/dynamic-ip-address-cyber-attacks>.
- [42] N. Sehatbakhsh, M. Alam, A. Nazari, A. Zajic, and M. Prvulovic. Syndrome: Spectral analysis for anomaly detection on medical iot and embedded devices. In *2018 IEEE international symposium on hardware oriented security and trust (HOST)*, pages 1–8. IEEE, 2018.
- [43] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18, 2017.
- [44] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart. Stealing machine learning models via prediction apis. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 601–618, Austin, TX, Aug. 2016. USENIX Association.
- [45] Wikipedia. Pipeline (computing). [https://en.wikipedia.org/wiki/Pipeline_\(computing\)](https://en.wikipedia.org/wiki/Pipeline_(computing)), 2004.
- [46] Wikipedia. Brute-force attack. https://en.wikipedia.org/wiki/Brute-force_attack, 2018.
- [47] Wikipedia. 20-year-old flaw found in ubiquiti networking gear running ancient php., 2020.
- [48] Wikipedia. *Brute-Force Attack*, 2020. https://en.wikipedia.org/wiki/Brute-force_attack.
- [49] Y. Xing, H. Shu, H. Zhao, D. Li, and L. Guo. Survey on Botnet Detection Techniques: Classification, Methods, and Evaluation. *Hindawi, Mathematical Problems in Engineering*, pages 1–24, Apr. 2021.
- [50] Q. Yang, P. Gasti, K. S. Balagani, Y. Li, and G. Zhou. USB side-channel attack on Tor. *Comput. Networks*, 2018.
- [51] Q. Yang, P. Gasti, G. Zhou, A. Farajidavar, and K. S. Balagani. On inferring browsing activity on smartphones via usb power analysis side-channel. *IEEE Transactions on Information Forensics and Security*, 12:1056–1066, 2017.
- [52] J. Yuan and S. Yu. Privacy preserving back-propagation neural network learning made practical with cloud computing. *IEEE Transactions on Parallel and Distributed Systems*, 25(1):212–221, Jan 2014.
- [53] Q. Zhang, C. Xin, and H. Wu. SecureTrain: An approximation-free and computationally efficient framework for privacy-preserved neural network training. *IEEE Transactions on Network Science and Engineering*, (in press), URL: <https://ieeexplore.ieee.org/document/9271910>.