Wireless Charging Power Side-Channel Attacks

Alexander S. La Cour*
Princeton University
Princeton, NJ, USA
lacour@princeton.edu

Khurram K. Afridi
Cornell University
Ithaca, NY, USA
afridi@cornell.edu

G. Edward Suh
Cornell University
Ithaca, NY, USA
suh@ece.cornell.edu

ABSTRACT

This paper demonstrates that today's wireless charging interface is vulnerable to power side-channel attacks; a smartphone that charges wirelessly leaks information about its activity to the wireless charger transmitter. We present a website fingerprinting attack and other preliminary attacks through the wireless charging side channel on iOS and Android devices. The website fingerprinting attack monitors the current draw of a wireless charger while the smartphone it charges loads a website from the Alexa top sites list. Our classifier identifies the website loaded on an iPhone 11 or a Google Pixel 4 with over 90% accuracy using wireless charging current traces. This attack represents a considerable security threat because wireless charging will always initiate when a compatible device is within the range of a charging transmitter. We find that the performance of the attack deteriorates as the contents of websites change over time. Additionally, this study finds that the wireless charging side channel is comparable to the wired USB charging side channel. Information leakage in both interfaces heavily depends on the battery level; minimal information leaks at low battery levels.

CCS CONCEPTS

• **Security and privacy → Side-channel analysis and countermeasures**.

KEYWORDS

wireless charging, side channel attacks, website fingerprinting

ACM Reference Format:
Alexander S. La Cour, Khurram K. Afridi, and G. Edward Suh. 2021. Wireless Charging Power Side-Channel Attacks. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21), November 15–19, 2021, Virtual Event, Republic of Korea.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3460120.3484733

1 INTRODUCTION

Smartphone charging has become increasingly prevalent. According to a Pew Research Center survey, 81% of American adults report owning a smartphone [36]. Moreover, a market research poll conducted by Veloxity, a phone charging station company, found that respondents charged their phones from 1.6 to 2.7 times per day [39]. While wired chargers are currently more common, the market share of wireless charging has been expanding. A BIS research report predicts the global wireless charging market will be worth over $20.97B in 2023, and the CEO of BIS Research has claimed that there will be more wireless chargers than charging cables by that time [5].

In this paper, we show that today's wireless charging interfaces are vulnerable to a power side-channel attack that can leak private information from a charging device to the transmitter of a wireless charger. In particular, we demonstrate the attack on the Qi standard [31], which is currently the dominant standard for wireless charging. The side-channel attack through wireless charging represents a substantial threat because it does not require a physical connection to a victim device and can occur without user permission or sophisticated equipment. While similar power side-channel attacks have been demonstrated through wired charging, wireless charging has been considered noisy and more secure against side-channel attacks. This paper is the first to investigate power side-channel attacks through wireless charging and demonstrate that practical attacks are feasible.

As a concrete example, we study a website fingerprinting attack through the wireless charging power side channel and perform detailed experimental studies on an Apple iPhone 11 and a Google Pixel 4. The phones are placed on a wireless charging transmitter and load a webpage from a set of candidates. As the webpage loads, we record the amount of current drawn by the wireless charging transmitter. After collecting enough data, a trained classifier can identify the webpage that corresponds to an unlabeled current trace. On 10-second duration current traces from both an iPhone 11 and a Google Pixel 4, we achieve an accuracy of over 90% and when the traces are truncated to 2.5 seconds we achieve an accuracy of at least 80%.

Our study also shows that this power side-channel attack does not rely on expensive or bulky measurement equipment such as a high-performance oscilloscope which makes concealing a power monitoring circuit in a wireless charger very plausible. In our experimental setup, we used a microcontroller to measure the current delivered to a wireless charger. We believe that when the charger is malicious, the adversary could place the attack circuitry inside the casing of the charger itself. Smartphone owners will generally not have access to the circuitry of public wireless chargers and will be unable to identify a malicious or compromised charger. Wireless public charging stations can be inserted in tables and chairs and are becoming ubiquitous [2]. There are currently over 200 smart devices that natively support the Qi standard [31], and older phones can implement the standard by connecting to a Qi-compatible wireless receiver via an accessory or case for as little as $10. Given the prevalence of wireless charging and the ease of an attack, we believe

*The work was done while the author was at Cornell University.

CCS '21, November 15–19, 2021, Virtual Event, Republic of Korea

ACM ISBN 978-1-4503-8454-4/21/11…$15.00
https://doi.org/10.1145/3460120.3484733

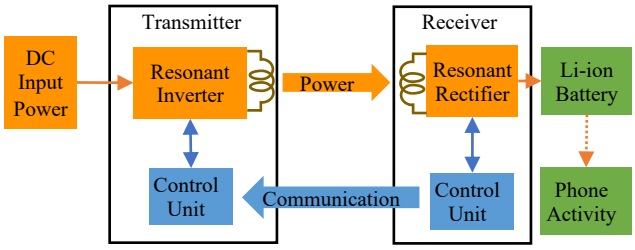

Figure 1: Transmitter and receiver hardware for the Qi standard.

that the side-channel attack through wireless charging represents a significant security risk.

In addition to demonstrating that today's wireless charging interface is vulnerable to practical power side-channel attacks, this paper also presents the results from a set of in-depth experimental studies to understand the capabilities and limitations of the wireless charging side channel. For example, we compare wireless charging and traditional wired charging in the context of power side-channel attacks. We find that the wireless charging side channel is comparable to the wired side channel in terms of classification accuracy despite some noise. We also observed the effects of other variables such as device type, the length of time between the collection of training and testing traces, and the trace length.

Our study also found that the amount of information leaked through these side channels in today's battery-powered devices depends heavily on the battery state of charge (SoC). When the battery SoC is high, the power consumption of the victim device is almost directly reflected on the power draw from the charger, revealing the activities on the device. On the other hand, when the battery SoC is low, most of the power from a charger is used to charge the battery. In that sense, we found that devices are far more vulnerable to wireless charging side-channel attacks when their battery level is above 80%. Unfortunately, given their convenience, users often leave devices on wireless chargers when fully charged. The chairman of the Wireless Power Consortium (WPC) stated that the WPC was unaware of any adverse consequences of prolonged wireless charging and suggested that topping off a phone battery will increase its life span [14]. For user privacy, our study suggests that future devices may want to adjust their charging algorithm and avoid fully charging a battery through an untrusted wireless charger.

Additionally, we also performed preliminary experiments on other potential side-channel attacks. The results suggest that the wireless charging power side channel can reveal information on the number of digits in a passcode, the number of white/bright pixels in an OLED display, the audio played while the screen is off, and the computations on a CPU.

The following summarizes the main technical contributions.

- This paper represents the first demonstration of the existence of a wireless charging power side channel on today's smartphones. Even with noise, this side channel leaks enough information to allow accurate website fingerprinting.
- This paper experimentally compares the wireless and wired charging side channels and shows that they leak the same power consumption information.
- This paper shows that the amount of information leaked through the charging side channel depends significantly on battery level.

2 BACKGROUND

This section provides technical background on wireless charging and power side-channel attacks which is necessary to understand the proposed wireless charging power side-channel attack.

2.1 Wireless Charging

The Qi open interface standard for wireless power transfer is the prevailing method for wirelessly charging smart devices. Qi was developed by the Wireless Power Consortium and describes the functional and physical characteristics necessary to allow the exchange of power and information between a receiver and a transmitter. Currently, Qi supports two power specifications to charge mobile devices: the Qi Baseline Power Profile, which delivers power below 5 W, and the Qi Extended Power Profile, which supports up to 15 W [41]. Wireless charging is becoming standard in new devices and since its release in 2008, Qi has already been integrated into over 200 smart devices [31].

Qi utilizes inductive charging to wirelessly transfer power from a transmitter to a receiver. Under this charging scheme, an inductive coil on the transmitter (the primary coil) couples to another coil on the receiver (the secondary coil). The transmitter runs an alternating current through its coil which induces an alternating voltage in the receiving coil by Faraday's law of induction. Additionally, capacitors connect to both inductive coils to form LC resonant circuits and enable resonant inductive coupling so that devices can charge even when up to 4 cm away [41]. The induced alternating voltage in the receiving coil is rectified and used to charge a battery or directly power a device.

Figure 1 shows the hardware implementation of the Qi standard, highlighting the electronics between the input power and the device battery. The communication between the transmitter and receiver occurs via backscatter modulation and is unidirectional from the receiver to the transmitter. The transmitting coil is powered by a resonant inverter while the receiving coil feeds a resonant rectifier. Both the transmitter and receiver contain communications and control units that actively regulate the power transferred to match the amount requested by the charging device.

The communication protocol of the Qi standard involves five phases. In the first phase, the power transmitter sends an analog ping to detect whether or not an object is present. The power transmitter then sends out a longer, digital ping to give the receiver time to reply with a signal-strength packet. If the transmitter considers this packet valid, it will continue to power its coil and proceed to the next phase. The third phase is known as the identification and configuration phase, where information is sent by the receiver in packets to properly configure the transmitter for power transfer. Next, the power transfer phase begins, during which the receiver sends control error packets to modify the supplied power. The final

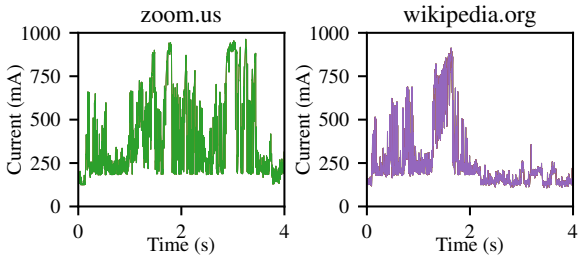

Figure 2: A wireless charger draws a varying amount of current as mobile webpages are loaded on the charging phone.

phase occurs when the receiver stops communication or requests the end of power transfer [41].

In terms of power delivery, Qi wireless charging is less efficient than wired charging. Wireless charging also introduces noise, and some have speculated that this type of noise is a good countermeasure against side-channel attacks that examine the amount of current used to charge a smartphone [22]. However, wireless charging transmitters do not store any significant amount of charge. Therefore, most of the current drawn by the transmitter will directly reflect the phone activity which acts as a load on the receiver.

2.2 Battery Charging Cycles

Most smartphones use lithium-ion (Li-ion) batteries. These batteries go through different charging stages [12]. The first stage, known as constant current charging, involves supplying the maximum allowable current to the battery, steadily increasing its voltage. Once the battery voltage reaches approximately 4.2 V, the second stage, known as constant voltage charging, will begin. During this phase, the supplied current drops off to limit the maximum voltage level of the battery. Once the battery SoC has reached 100%, the charger will provide a topping charge to make up for any discharging and return the SoC to 100% [40].

As a result of these charging stages, the amount of current drawn by a phone from a charger heavily depends on the battery SoC and may not be affected by how much power the phone consumes. For example, when a phone's battery is at a low state of charge, corresponding to constant current charging, the amount of power the phone consumes will not significantly affect its overall current draw. This is because the current draw is already at its maximum without the phone consuming any power. Any power consumed by the phone reduces the current charging the battery but does not affect the current draw. On the other hand, during constant voltage charging, the power consumption of the phone will reduce the battery voltage and a larger current will be drawn to offset this. When the phone battery is fully charged, no current flows into the battery and the amount of power drawn from the charger is a direct reflection of the power consumed by the phone.

2.3 Power Side-Channel Attacks

Side-channel attacks are methods to acquire sensitive information through unintended secret-dependent variations in physical behaviors. The information leaked from a side-channel attack is a byproduct of operations occurring on hardware and is not a specific software vulnerability.

Power side-channel attacks are a specific type of side-channel attack that analyze the power traces of the electrical activity on a device to extract information [21]. Simple power analysis (SPA) is a method of power side-channel attack that infers a secret value from a power trace by identifying power consumption profiles that directly depend on the secret. Frequency filters and averaging functions are sometimes applied to filter out noise in these power traces [9]. Differential power analysis is a more complex side-channel attack that allows the identification of intermediate values within cryptographic computations after a statistical analysis of prior collected data.

While power side-channel attacks are an established field of research, applying these techniques to mobile devices is a relatively new endeavor. Mobile devices are uniquely susceptible to side-channel attacks because they are portable, continuously powered on, and have many sensors. Understanding the extent of sensitive information that a power side-channel attack can infer will provide insight into security risks.

Smartphone security relies on two premises: application sandboxing and a permission system. These ensure that applications cannot access sensitive information contained in another resource. Yet, even without direct access to the data pins of a smart device, power side-channel attacks have proven to be effective. For example, Yang et al. [44] showed that charging a smartphone over a USB cable exposes a side channel that is vulnerable to an SPA attack. By monitoring the power that a charging smartphone drew, they successfully inferred private browsing information. Figure 2 shows that in the current traces we collected, different websites leave unique signatures through the wireless charging side channel over short time durations.

3 POWER SIDE CHANNELS IN WIRELESS CHARGING

This section introduces the concept of wireless charging power side-channel attacks and discusses their capabilities and limitations at a high level. The following section provides a more in-depth study using website fingerprinting as a concrete example attack.

3.1 Threat Model

Figure 3 shows the threat model that is assumed for the wireless charging side-channel attack. Under this threat model, an attacker can monitor and record the amount of power delivered to an untampered Qi wireless transmitter from a compromised public wireless charging station. The target device performs activities that depend on sensitive events or data values, influencing its power consumption. The attacker's goal is to infer the events or data values on the target device by analyzing the recorded power traces. While we assume the public charging station is compromised, it need not be malicious because the classification and inference can occur remotely.

Wireless charging does not require user permissions or initiation and will begin if both the mobile device and the transmitter follow the Qi standard and are in range (4 cm). There is no need for the device to plug into the charging station. The target device is not

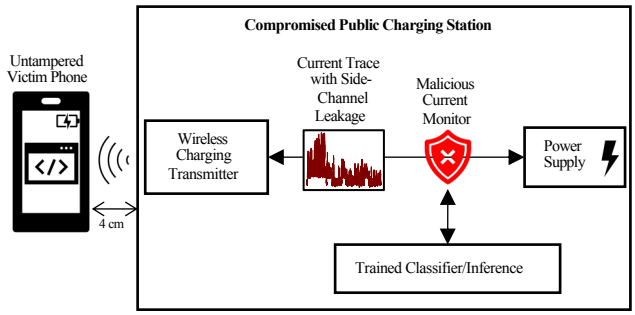

Figure 3: Threat model demonstrating a power side-channel attack by a compromised public charging station.

assumed to have any malicious software and this threat model does not depend on any particular software vulnerability. Additionally, this type of attack does not require any physical tampering of the target device or battery.

3.2 Experimental Setup

The high-level idea of the wireless power side-channel attack is similar to that of the traditional wired power side-channel attack. However, wireless charging interfaces do not have physical wire connections and are likely more susceptible to noise. In that sense, the main technical contributions of this paper lie in experimental studies that demonstrate that wireless power side-channel attacks are feasible in the mobile phones of today and their capabilities are comparable to those of wired power side-channel attacks.

Here we briefly describe the experimental setup that we used. The experiments are designed to understand the capabilities and limitations of the wireless power side channels:

- Does the wireless power side channel leak enough information to infer activities on a mobile device even with noise in the wireless interface? Are the measurements repeatable?
- How is the wireless power side channel impacted by the battery level?
- How does the wireless power side channel compare to the wired power side channel in terms of leakage?

Current Trace Collection Circuit. The DC delivered to either a 5 W Adafruit Qi Wireless Charging Transmitter or a 10 W Max Anker Wireless Charging Pad from a USB AC adapter was sampled by placing an INA219 High Side DC Current Sensor in series with the V_{CC} wire of the Micro-USB cable that powered the transmitters. This is depicted in Figure 4. An Arduino Micro sampled the current sensor at a frequency of 700 Hz (500 Hz in Sections 5.6-5.8). The cost of the entire current trace collection circuit used in this work is less than $30.

Example Current Traces. Figure 5 demonstrates that like the USB charging side channel, the wireless charging side channel also leaks enough information to distinguish different websites. Additionally, we find that the collected current traces are repeatable across different trials indicating that the activity visible in the traces is a direct result of loading a particular website. In all cases, the websites take

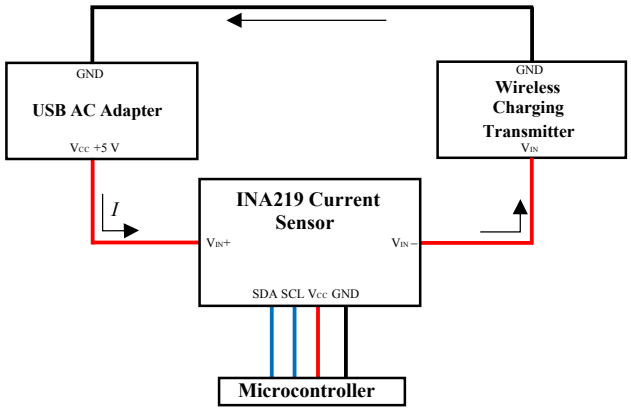

(a) Overview of current trace monitoring.

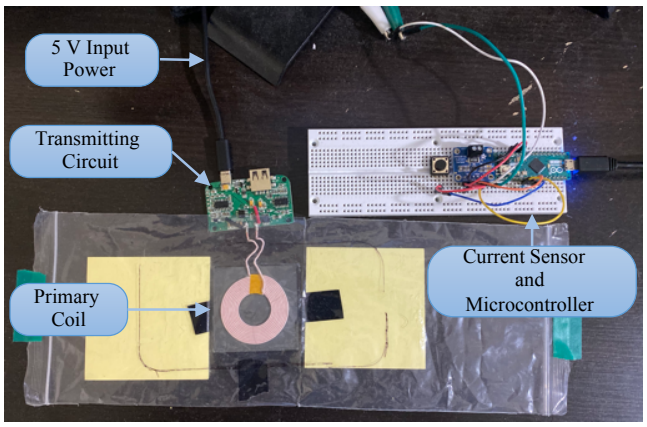

(b) Photo of setup with the Adafruit 5 W transmitter.

Figure 4: The current trace collection setup used in all experiments.

a variable amount of time to load, and once fully loaded, the current drawn by the charging transmitter returns to a steady level.

Phone Configuration. The attack is performed on an Apple iPhone 11 (2019) running iOS 14 and a Google Pixel 4 (2019) running Android 11 which are both capable of wireless charging with Qi-certified chargers up to powers of 7.5 W and 11 W respectively. When the iPhone 11 traces were collected without noise, an outline for the phone was placed around the coil so that it could be positioned consistently above the transmitter across every trace. Otherwise, both phones were placed at various orientations while remaining centered enough to properly charge.

3.3 Impact of Battery Level

Figure 6 shows how the wireless charger's current draw varies as the charging phone's battery level increases. The results indicate that the charging profiles of a wireless charger mirror those of a wired charger [18]. At a low SoC, the current draw is relatively fixed except for a high-frequency component coming from the wireless interface. Then, the power draw gradually decreases as the battery state of charge increases.

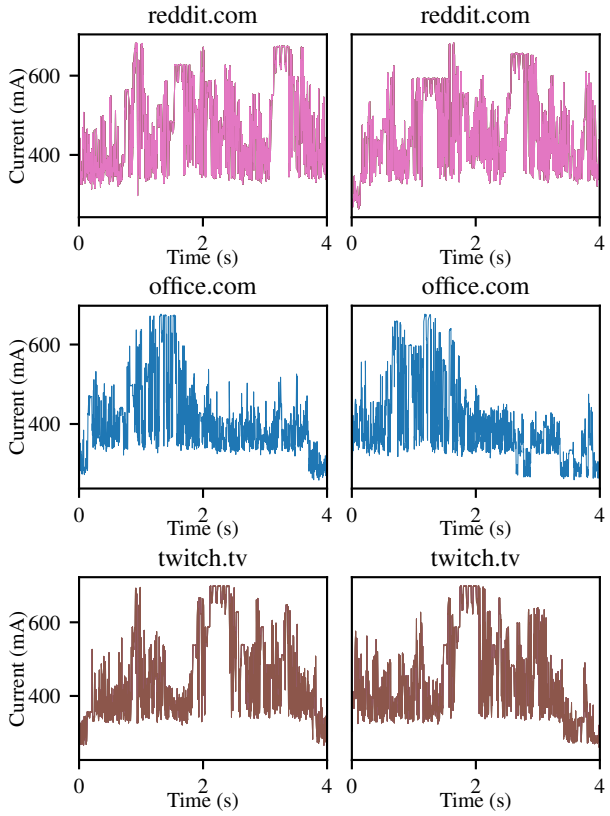

Figure 5: Current traces demonstrating the activity leaked when automatically loading webpages on an iPhone 11.

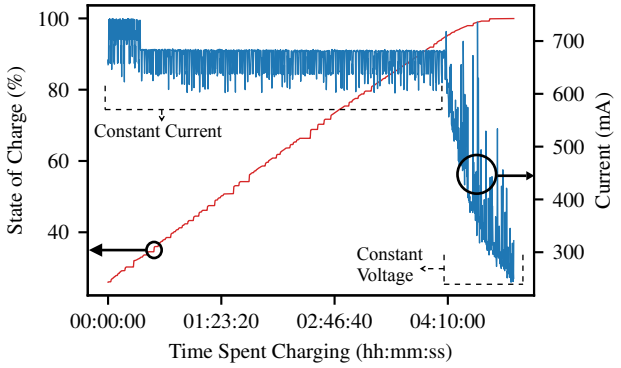

Figure 6: Current delivered by a 5 W Qi charger/battery state of charge vs charging time for an iPhone 11. The constant current and constant voltage charging stages are identified.

Figure 7 shows how the average current consumption of a wirelessly charging phone varies as it executes different processes. The experiment was carried out at 8 different battery levels. While the results demonstrate that different processes consume different amounts of power on average while wirelessly charging, a clear differentiation between activities only occurred when the SoC was high. When the state of charge is less than or equal to 95%, the activities were generally indistinguishable by the metric of average

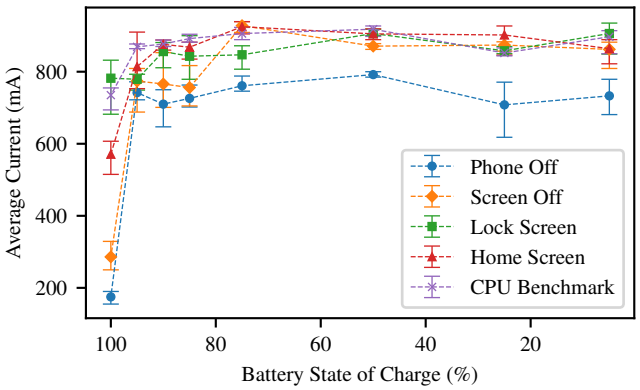

Figure 7: The average current consumption vs iPhone 11 state of charge for five different activities.

current assumption. The reason for this is that when the phone's battery is fully charged, the amount of power delivered by the wireless transmitter is solely determined by the power the phone is currently using as it cannot deliver more charge to a battery that is already at maximum capacity. If the battery is not fully charged, the power consumption of an app running on the phone may not dominate the power draw from the charger, as much of it will be used to charge the battery.

Even if the average power consumption does not leak enough information to distinguish different activities at a lower battery level, a trace of dynamic power consumption over time can reveal far more information. For all experiments in our evaluation section, except for Section 5.8 where different battery levels were examined, current traces were collected automatically beginning when the device's battery was full. During the duration traces were collected, the device's state of charge fell but always remained above 90%. In general, we found that battery-powered mobile devices are more susceptible to power side-channel attacks when the battery state of charge is high. The exact amount of information leaked depends on the charging algorithms used by a victim device. Our experiments in Section 5 suggest that even with time-series data, the iPhone 11 leaks little information when the battery charge level is below 80%.

4 WEBSITE FINGERPRINTING ATTACK

4.1 Attack Overview

The attacker seeks to utilize collected power data to identify the webpages being loaded in a mobile browsing application by a victim as they wirelessly charge their phone. As established by the mobile power side-channel attacks previously discussed, loading a website on a smartphone can affect its power consumption patterns. When the phone battery is near full charge, the power delivered to the wireless charging transmitter is directly proportional to the fluctuations in activity on the phone and can be recorded by a compromised public wireless charging station.

A set of training data can be collected by the repeated loading of websites onto a charging device in this manner. This data can then be preprocessed and fed to a website fingerprinting classifier for training and validation. After training, the model can classify new power data collected from victims by compromised charging

Table 1: 1D CNN model architecture where the duration of the windowed input trace is 1 second (699 samples). This window is split into three slices so that the LSTM layer can learn the chronological relationship between the features from each slice.

Layer	Type	Output Shape
0	Input Layer	(3,233,1)
1	TimeDistributed(Conv1D)	(3,229,128)
2	TimeDistributed(MaxPooling1D)	(3,114,128)
3	TimeDistributed(Conv1D)	(3,110,192)
4	TimeDistributed(MaxPooling1D)	(3,55,192)
5	TimeDistributed(Conv1D)	(3,51,300)
6	TimeDistributed(MaxPooling1D)	(3,25,300)
7	TimeDistributed(Flatten)	(3,7500)
8	LSTM Layer	(900)
9	Dropout Layer	(900)
10	Fully-Connected Layer	(900)
11	Fully-Connected Layer	(50)

stations. This victim data will then be similarly preprocessed to form the testing data, which if classified correctly, will reveal an individual's private browsing activity. This attack is performed on untampered wireless charging transmitters, but a malicious transmitter designed for these attacks could provide more accurate traces.

4.2 Current Trace Collection

In the case of the iPhone 11, the mobile Safari browser connects to the Safari development tool, Web Inspector, on a Mac computer. The computer then runs a script that sequentially loads a set of websites on the iPhone up to 50 times. We collect separate data for the wireless and wired chargers. Trace collection on the Pixel 4 followed a similar process except that the Chrome browser and Chrome Developer Tools were used to initiate webpage loading. The current trace corresponding to the first 10 seconds of loading a website is recorded and between loading each site, the script waits 4 seconds. This script also automatically initializes the data collection to ensure that all power traces are synchronous and aligned. The top 20 and 50 non-adult websites from the Alexa Top Sites in United States list [17] were examined in this attack. All websites we visit in this experiment are listed in Table 7 and utilize a secure connection via HTTPS which is encrypted with TLS.

For nearly all configurations, we collect testing traces with the intent to mirror standard device operation. This included setting the phone's brightness and volume at a constant level (although no websites visited automatically played audio) and enabling Bluetooth and cellular data. The exception to this is in Section 5.6, where we collect test traces with volume, Bluetooth, and cellular data disabled. For all traces, notifications on the devices were disabled to prevent calls from interrupting the data collection script. The Pixel 4 did not have a SIM card inserted, so it did not have cellular data enabled.

4.3 Classification Algorithm

For feature extraction, we broke each current trace into segments that represented 1 second of the original trace, with 97.5% overlap. These segments were acquired by applying a sliding window algorithm to the overall current trace. We selected this feature duration because many of the identifiable features that distinguished each trace were less than a second long. Training on many small segments rather than entire traces helped to increase the amount of training data available and reduced overfitting by making our model more shift-invariant. Each trace in the test set is broken into segments as was done for the training data, and each segment's classification is cast as a vote for classifying the overall test trace. The final trace label was assigned using a majority voting scheme. A 64/16/20 training/validation/testing split was used.

Deep neural networks act as both feature extractors and classifiers, which can make attacks more successful than traditional techniques. Additionally, convolutional neural networks (CNNs) [20, 34, 45] incorporate translation invariance, which allows them to recognize features even if they are translated to different time positions. Although our current traces were collected automatically, the loading time of pages sometimes is delayed randomly due to website traffic or other causes.

A 1D CNN, the architecture of which is detailed in Table 1, is trained as a classifier on these segments and was implemented in the Keras [8] software package. Our architecture is a modified version of a 1D CNN that was used for human activity recognition [7]. This model was chosen as a base because it was designed for multi-output classification, had a foundational architecture that was easy to build on, and proved resilient to overfitting.

The topology of our CNN is three convolutional layers followed by a long short-term memory (LSTM) layer [35], a fully connected layer, and a Softmax layer with one output for each website. Every convolutional layer used ReLU activation [1], had a convolutional window of size 5, and was followed by a max-pooling layer with a window of size 2 and a stride of 2. Each window was split into three equal-length temporal slices to allow the LSTM layer to update its weights based on the chronological relationship it learned between the features from each slice. The CNN layers were wrapped in a TimeDistributed layer which is a layer that applies the same input operation across all time slices constructed from each window. There are 128 filters in the first convolutional layer, 192 in the second, and 300 in the third. The network also uses a dropout layer with a frequency of 50% to further reduce overfitting by randomly dropping nodes and regularizing the network.

We chose the LSTM layer for this classification problem because it is a recurrent neural network layer that can learn the order dependence within data. Given that the segments the network examines are 250-350 time steps in length, the ability of the classifier to learn order dependence would allow it to identify the presence of multiple features within a single segment. The data we collected was a one-dimensional time series and while loading a website, many events such as executing JavaScript and loading images will always be executed by the phone in the same order. In this way, the LSTM layer complements the convolutional layers in our architecture: the convolutional layers extract features and the LSTM layer learns their order dependence.

Table 2: Rank-1 and rank-2 accuracy (%) for 1D CNN model when classifying 20 websites with a fully charged iPhone 11.

Current Trace Type	10 s	6 s	5 s	4 s	2.5 s
Noiseless Wireless Rank-1	94.0	94.5	94.0	87.5	80.5
Wireless Rank-1	N/A	87.0	87.5	87.5	82.0
Noiseless Wired Rank-1	97.0	96.0	96.5	96.0	88.5
Noiseless Wireless Rank-2	96.0	96.5	97.5	94.0	88.0
Wireless Rank-2	N/A	94.0	94.0	89.5	87.0
Noiseless Wired Rank-2	99.0	97.5	98.0	97.0	93.5

The CNN outperformed all other classifiers we explored when evaluated on our collected data. The second-best performance we obtained was with a Random Forest [6] classifier that was trained with the frequency domain representation of the current traces. Although we were able to get reasonably high accuracy with this classifier, it was not able to generalize well to different charging conditions. In contrast, our CNN performed well on all scenarios in which current traces were collected and did not require any feature engineering aside from the application of the sliding window algorithm. Our classifier also successfully identified traces that were time-shifted with respect to the training data. Overall, our attack can effectively classify traces collected from multiple devices and charging methods with the same feature extraction process. This is critical because our threat model is intended to apply to a variety of phone models, operating systems, and chargers.

5 EVALUATION

In this section, we present our findings and detailed experimental results on the website fingerprint attack through wireless charging. Rank-1 and rank-2 identification accuracy of the classifier in different scenarios were calculated. Rank-1 counts a classification as correct if the majority vote picks the correct website for the trace. Rank-2 accuracy counts a classification as correct if either the website with the most or second-most votes is correct. The baseline accuracy of a random guess classifier for the 20 websites is 5% for rank-1 and 10% for rank-2.

We conducted a range of experiments aiming to identify how the classifier accuracy changed with respect to the following variables: (1) device manufacturer; (2) number of websites visited; (3) different devices for training and testing; (4) different chargers methods for training and testing; (5) length of current traces; (6) noise; (7) aging of training traces; (8) battery state of charge. The following subsections detail our findings and contributions concerning each question.

5.1 iPhone 11 vs Pixel 4

In this subsection, we aim to identify how the accuracy of the classifier depends on the device used to collect current traces. The iPhone 11 and Google Pixel 4 were both used to collect current traces under a variety of conditions. Results from these experiments are reported in Table 2 (iPhone) and Table 3 (Pixel). All test traces in this section, unless otherwise specified, included noise in the form of normal device operation conditions such as leaving the phones' Bluetooth, cellular data, volume, and notifications on while placing them at a variety of alignments with the transmitting coil.

Table 3: Rank-1 and rank-2 accuracy (%) for 1D CNN model when classifying 20 websites with a fully charged Pixel 4. All traces were collected under normal operation conditions.

Current Trace Types	6 s	5 s	4 s	2.5 s
Wireless Rank-1	95.0	94.0	95.5	85.5
Wired Rank-1	74.0	75.0	70.5	63.0
Wireless Rank-2	97.5	98.0	96.5	91.5
Wired Rank-2	83.0	85.5	82.5	79.0

Table 4: Rank-1 and rank-2 accuracy (%) for 1D CNN model when classifying 50 websites with a fully charged Pixel 4. All traces were collected without noise.

Current Trace Types	5 s	4s	3 s	2 s	1.5 s
Wireless Rank-1	98.8	98.4	98.0	96.8	92.8
Wireless Rank-2	99.2	98.8	98.4	97.6	96.8

The classifier achieved a rank-1 accuracy of at least 82.0% and a rank-2 accuracy of at least 87.0% when classifying wireless traces from the iPhone 11 with trace durations ranging from 2.5 to 6 seconds. Pixel 4 wireless traces were classified with higher accuracy, especially at longer trace lengths. It achieved a rank-1 accuracy of at least 85.5% and a rank-2 accuracy of at least 91.5% with trace durations ranging from 2.5 to 6 seconds. The high accuracy of the classifier in these scenarios indicates that the small changes in phone activity that occur while loading various websites are detectable through this wireless side channel in both devices examined.

5.2 Extended Website Set

In addition to the website fingerprinting attack on the Alexa top 20 sites we also experimented with a larger data set consisting of current traces from the Alexa top 50 sites to further demonstrate the capabilities of this attack. The same current trace collection setup from Section 3.2 was implemented, although only 25 traces were collected for each website, the traces were collected without noise, and the maximum length of each current trace collected per site was 5 seconds instead of 10. Additionally, the classifier's final fully-connected layer was adjusted to fit 50 classes instead of 20. We found that the classifier can identify the current traces with a rank-1 accuracy of 98.8% when the traces were five seconds long (see Table 4).

5.3 Training and Testing on Different Devices

In order to see whether or not a cross-device attack is possible in this threat model, we trained the classifier exclusively on current traces from the iPhone and tested on traces from the Pixel and vice versa. When the current traces from a different device were used for training, the classifier was unable to identify traces from the device at all. Training on iPhone traces and testing on Pixel traces resulted in a rank-1 accuracy of 4.2% which is worse than a random

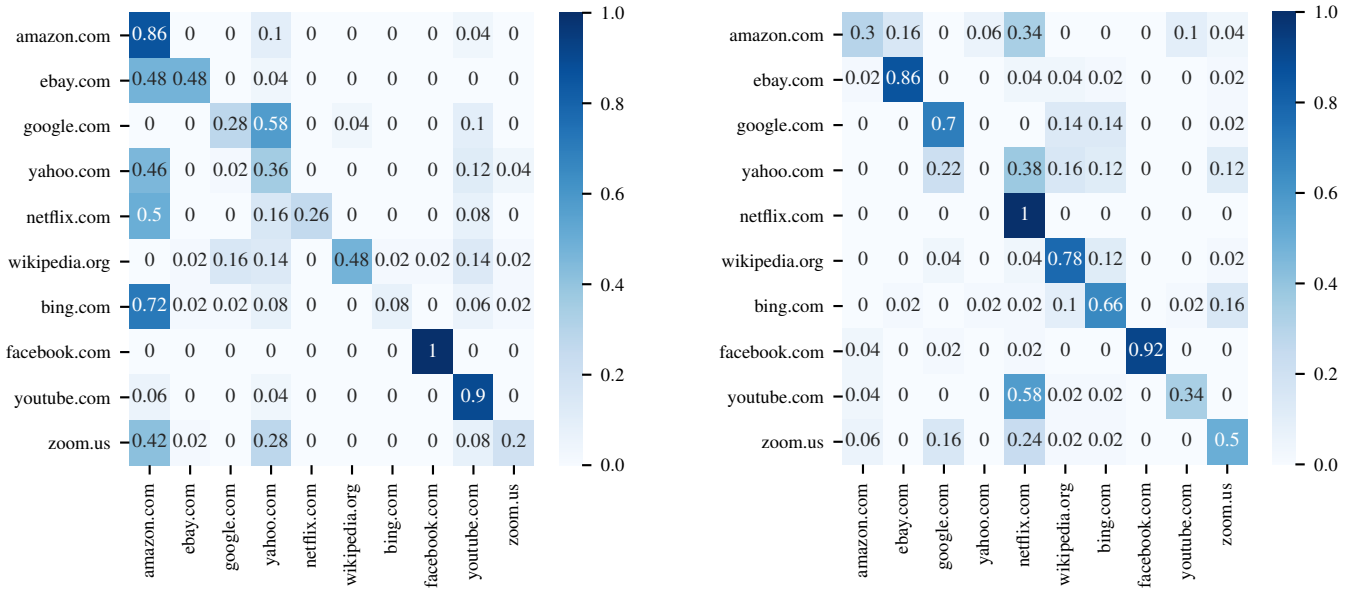

(a) Training on wireless traces, testing on wired traces.

(b) Training on wired traces, testing on wireless traces.

Figure 8: Results from training and testing across different chargers on traces collected with the iPhone 11. The vertical axis shows the true label and the horizontal axis shows the predicted label. An ideal classifier would have ones down the diagonal.

guess and a rank-2 accuracy of 12.1% which is only slightly higher than that of a random guess. Training on Pixel traces and testing on iPhone traces was no better. In this scenario, the classifier achieved a rank-1 accuracy of 5.7% and a rank-2 accuracy of 11.6%.

These results align with the findings from previous studies that found a drop in classification accuracy resulting from training and testing on different devices. This indicates that the power consumption variations depend on individual device characteristics and that the information extracted by the classifier from current traces depends on the charging device. An effective realistic attack would likely need to train on traces from a variety of phones to be able to generalize and account for more trace variety.

5.4 Training and Testing on Traces from Different Chargers

Current traces from a wired, 5 W charger were also collected with both the Pixel and the iPhone. Unlike wireless traces, wired traces from the iPhone were classified with higher accuracy than those of the Pixel. The minimum rank-1 and rank-2 accuracies of the classifier on the wired iPhone traces were 88.5% and 93.5%, respectively, whereas they were 63.0% and 79.0% on the Pixel.

Across all device and charger combinations, our classifier was able to perform well without any preprocessing or changes to the architecture. The accuracies achieved by the classifier when trained and tested on wired and wireless traces are similar, indicating that the information leakage from the wireless charging power side channel is comparable to that of the wired charging power side channel for the same device. In the case of the Pixel 4, the wireless current traces were identified with higher accuracy than the wired current traces.

Figure 10(a) shows the current traces measured using wired and wireless chargers while loading zoom.us on iPhone 11. A visual comparison suggests that the wired and wireless channels leak the same information when a website is loading; the patterns in the current traces when the phone is fully charged are similar. The traces differ in that the wireless traces contain a signal with a frequency of approximately 11 Hz and appear to be noisier than the wired traces.

In order to measure how comparable both charging side channels are, the classifier was trained exclusively on current traces from the wireless charger and tested on traces from the wired charger and vice versa. Using 10 websites and 2.5 second long traces, the classifier identified websites from the iPhone correctly with significant accuracy. The results of this experiment are shown in Figure 8. Training on wired traces and testing on wireless traces produced a rank-1 accuracy of 60.6% compared to a baseline of 10% and a rank-2 accuracy of 75.0% compared to a baseline of 20%. Training on wireless traces and testing on wired traces achieved a rank-1 accuracy of 49.0% and a rank-2 accuracy of 68.4%. The only website that was identified with over 90% accuracy in both situations was facebook.com.

The existence of cross-channel leakage across both wired and wireless charging indicates that wirelessly charging devices may be susceptible to existing USB power side-channel attacks that have been trained only on wired power data.

5.5 Impact of Trace Duration

In addition to the full duration traces, the classifier was trained and tested on shorter duration traces. These were formed by taking a slice of the first n seconds of data from the original trace. Out of all trace lengths examined, the best wireless and wired rank-1 identification accuracies achieved were with 5-second traces and 6-second traces respectively. While the classifier performed the worst on 2.5-second traces, the overall identification accuracy was still quite high and close to the best rank-1 accuracies out of all trace durations. These shorter traces removed noise present in the full 10-second traces because the websites examined take approximately 4 seconds to load [33]. However, most websites take over 2.5 seconds to load, so traces of this duration cut off part of the signal from the website loading and therefore deteriorated identification accuracy. Furthermore, websites that autoplay videos had consistent leakage in their traces even after they initially loaded.

5.6 Impact of Noise

As evidenced by the results discussed in Section 5.1, the attack is quite resilient to noise and was able to identify the test traces with high accuracy, even though the circumstances of the device varied between training and testing traces. This demonstrates that our attack is feasible in realistic scenarios where the current trace collected while a website is loading may be corrupted or altered by the existence of other executing processes.

In order to measure how well the attack might perform without noise, current traces were collected from the iPhone 11 while the volume, Bluetooth, and cellular data were disabled at a sampling frequency of 500 Hz. Additionally, an outline from the phone was placed over the charger so that the alignment and angle of the phone over the transmitting coil were consistent.

The classifier performed slightly better when trained and tested on the noiseless traces compared to those collected under normal operating conditions. The full results are reported in Table 2. When classifying noiseless wireless traces, the classifier obtained a rank-1 accuracy of at least 80.5% and a rank-2 accuracy of at least 88.5% with trace durations ranging from 2.5 to 10 seconds. We present the confusion matrix for 5-second traces in Figure 9. For comparison, noiseless wired traces collected under the same conditions resulted in a rank-1 accuracy of at least 88.5% and a rank-2 accuracy of at least 93.5%.

5.7 Impact of Length of Time Between Trace Collection and Testing

In this scenario, training and testing traces were collected on the same iPhone 11 except the test traces were collected nine months after the training traces were collected. Table 5 summarizes the results of this scenario. Many of the websites we studied had dynamic content, such as news. After many months, the media in these websites completely changed which resulted in the current traces altering as well. Although accuracy was significantly lowered in this experiment, the classifier still performed over four times better than a random guess would achieve at some trace lengths.

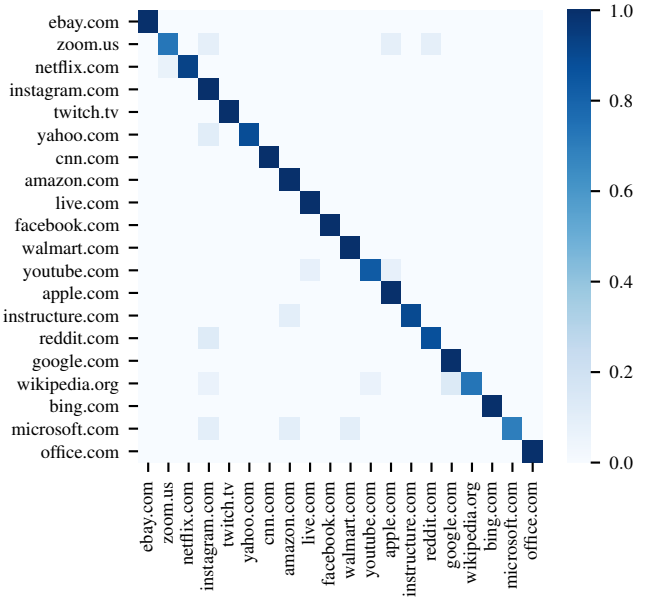

Figure 9: Confusion matrix for the classification of 200 unlabeled 5-second current traces across 20 websites collected on the iPhone 11 without noise.

Table 5: Rank-1 and rank-2 accuracy (%) for 1D CNN model when classifying with old training data.

Current Trace Type	6 s	5 s	4 s	2.5 s
New Traces Rank-1	18.0	20.5	22.5	13.5

5.8 Impact of Battery State of Charge

Below approximately 80% state of charge, both wired and wireless charging side channels in our experiments do not leak enough information for the classifier to identify the traces with any significant accuracy. This is shown in Figure 10. For the wired channel, information begins to be revealed when the battery state of charge reaches approximately 95%. The wireless channel could consistently classify traces with a battery state of charge as low as 90%.

Figure 10 also reveals how the power side channel through wired charging is affected by the battery level. The variations from the phone's activities are visible at higher battery levels but not at lower ones.

Previously, Yang et al.[44] found that power traces collected at battery levels of 30% were classified with accuracy almost as high as those collected when the battery was fully charged. The discrepancy seems to suggest that the newer smartphones are more resilient to power side channels. In order to further investigate how this side channel is affected by battery levels, current traces were collected from older Apple iPhone models, an iPhone 6s, and an iPhone 8, and compared on the same scale. We collect the power traces for both of these phones using the same data acquisition setup that we used with the iPhone 11.

Wired traces collected on an iPhone 6s leaked activity at lower battery levels than the iPhone 11 did. This can be seen in Figure 11.

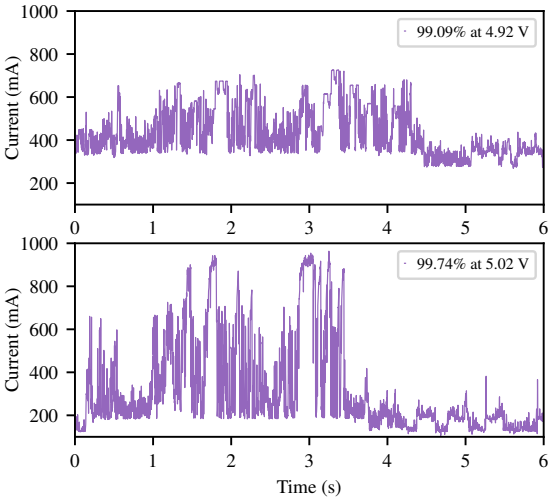

(a) wirelessly charging (top) and wired charging (bottom)

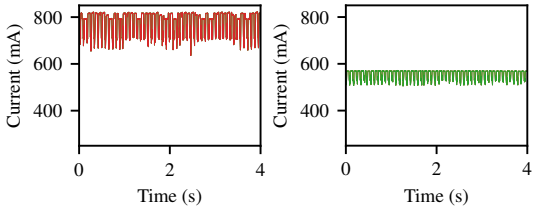

(b) 76.22% at 5.10 V

(c) 50.21% at 4.87 V

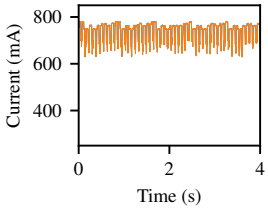

(d) 32.59% at 4.81 V

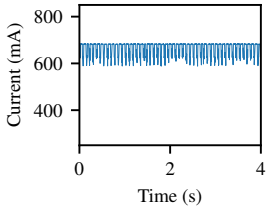

(e) 21.09% at 4.58 V

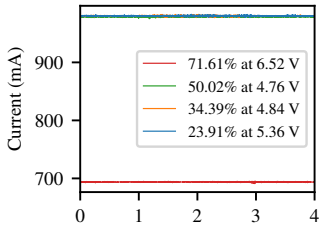

(f) While USB charging, the current traces recorded at the lowest three battery levels are indistinguishable.

Figure 10: The current traces for wireless and wired charging when loading zoom.us on an iPhone 11 for different battery levels (SoC and voltage). Plots (b)-(e) depict wireless charging.

The activity was visible at battery levels as low as 50% but became obfuscated at battery levels of 30% or lower. The iPhone 8 wired power side channel revealed activity in the same range of battery state of charge as the iPhone 6s.

While the iPhone 6s does not support wireless charging, the iPhone 8 is Qi-compatible. Its wireless current traces do not leak any significant information when the battery level is less than or equal to 70%. It is possible that there was a change in the hardware design between the iPhone 8 and iPhone 11 that removed the USB power side channel at battery levels below full charge. However, even on the iPhone 8, little activity was revealed at the 30% battery level compared to the Android phones studied by [44]. Additionally, even though the iPhone 11 is not as vulnerable to USB power side-channel attacks as the iPhone 8, both phones appear to be similarly susceptible to the wireless charging side-channel attack at higher battery levels.

6 OTHER ATTACK EXAMPLES

This paper demonstrates the website fingerprinting attack as an example of a wireless power side-channel attack. However, the wireless power side channel has the potential to leak other types of information about activities on a mobile device that affect the device's power consumption. The wireless charging interface may also introduce additional vulnerabilities beyond side-channel information leakage. For example, a malicious wireless charger may deliver a high current as a way to damage a circuit or perform repeated charging/discharging cycles to reduce battery life.

In this section, we discuss other side-channel attack examples through the wireless charging power side channel with preliminary experimental results that show their feasibility.

6.1 Estimating Passcode Length

The power consumption of a mobile device is sensitive to touch screen inputs. We found that the wireless power side channel can leak information about a user's passcode. Current traces collected during a passcode input show a momentary increase in the current consumption for each digit of a passcode entered by tapping the screen. Although the current surges do not directly reveal individual digits, the length of a passcode can be visually discerned from a current trace collected during its input by counting the number of surges it contains. Figure 12 reflects the entry of three different passcodes of varying lengths. The entry of a digit results in significantly more power being drawn than when the phone is at rest and the length of each passcode can easily be identified.

Knowing a passcode's length significantly reduces the search space needed to crack it, especially when combined with other information extraction attacks such as smudging attacks.

6.2 OLED Screen Power Consumption

The Google Pixel 4 in our experiments uses an OLED display, which is a type of display in which light is emitted by individual diodes and not by a backlight. This type of display generally consumes less power because individual pixels only light up if required by the screen content.

Using the wireless charging power side channel, we found that the current draw of the phone directly correlates with the number of white pixels on the screen. To measure the power consumption as a function of the number of white pixels, a completely black

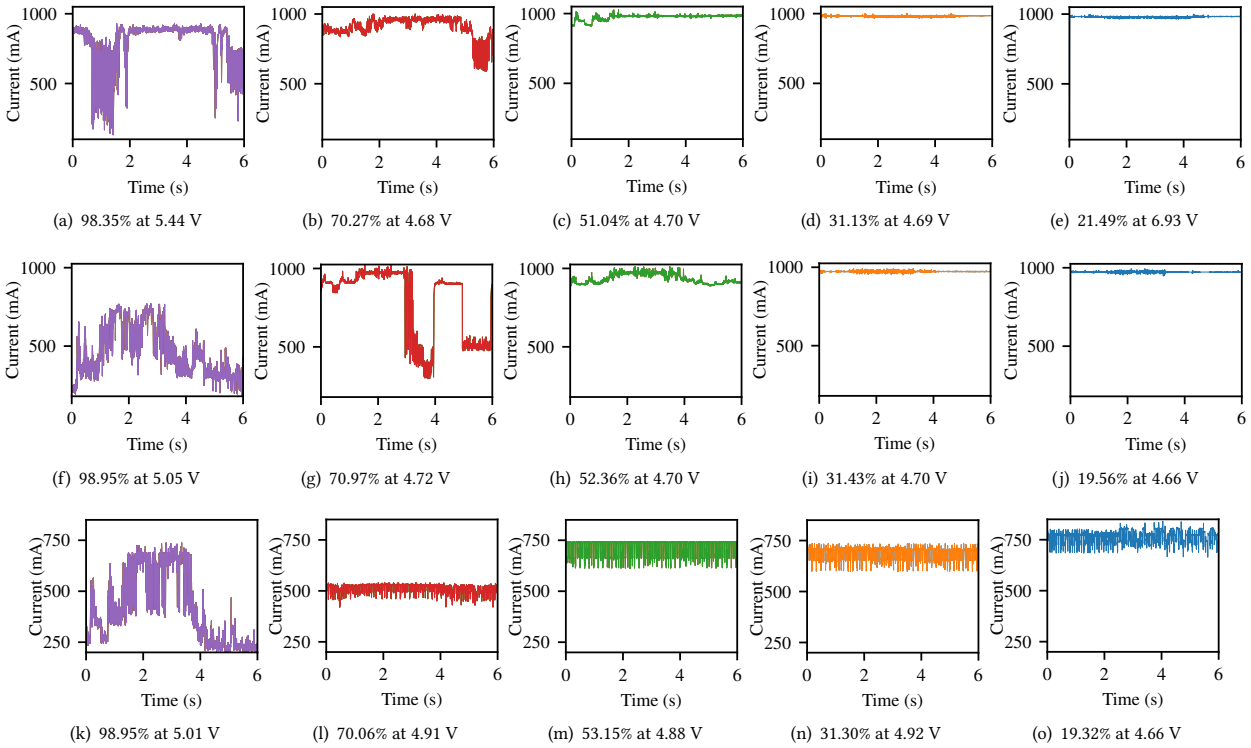

Figure 11: Current traces for loading zoom.us on different devices while wirelessly charging and USB charging: wired iPhone 6s (top), wired iPhone 8 (middle), wireless iPhone 8 (bottom).

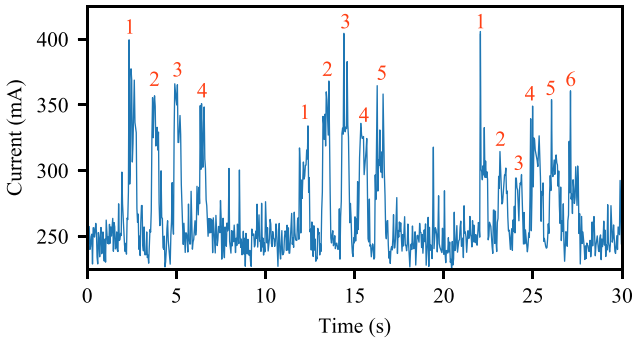

Figure 12: The wireless charging current trace reflecting the input of three passcodes of lengths of 4, 5, and 6 digits. A few seconds of inactivity separate the entry of each passcode. Each digit corresponds to a momentary surge in current draw.

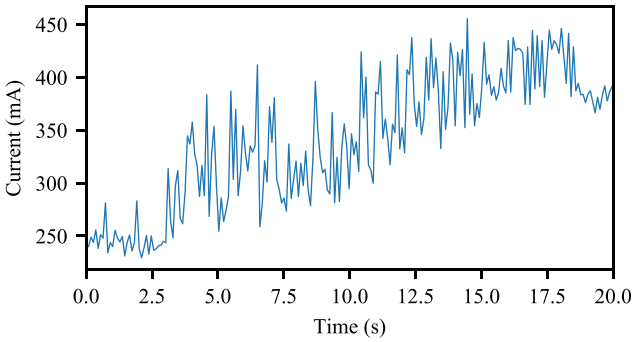

Figure 13: The current consumed by the wireless charger increases linearly on average as more pixels become white.

image is displayed on the screen, and a white image is slid across it, increasing the percentage of white pixels on the screen. Figure 13 displays the current trace collected during this process where the average current draw linearly increases as a function of the number of white pixels on an otherwise black screen.

With more precise power measurement devices, this side channel has the potential to leak even more information on the screen content, such as the type of notifications on a screen, or whether there is an incoming call.

6.3 Audio Fingerprinting

The previous subsections show that the wireless power side channel can leak user interface (UI) activities, especially from the variation in the screen power consumption. Here, we use an audio fingerprinting attack to demonstrate that the wireless power side channel can also leak background activities.

Table 6: Rank-1 and rank-2 accuracy (%) for 1D CNN model when classifying audio files played on the Pixel 4 speakers.

Current Trace Type	4 s	3 s	2 s	1.5 s
New Traces Rank-1	88.0	88.0	84.0	82.0
New Traces Rank-2	90.0	92.0	88.0	84.0

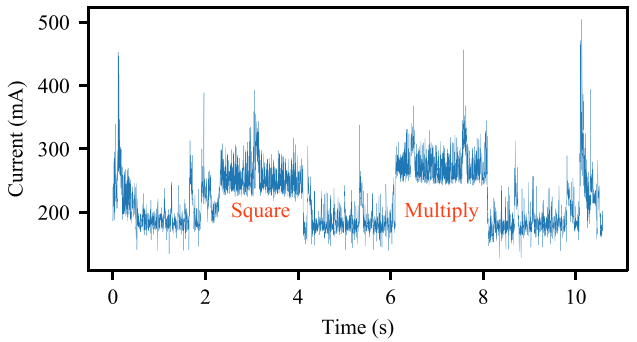

Figure 14: The current trace that demonstrates distinction between idle time and repeated square/multiply operations.

In this attack, we use the classification algorithm from the website fingerprinting attack to perform a fingerprinting attack on an audio track that was playing on the Pixel 4. With the phone screen off, we automated the phone to play 10 locally stored audio files from the LJ Speech Dataset [19] in a round-robin fashion. We then formed training, validation, and testing data sets that were used to determine which audio track was playing at the time each current trace was recorded. The full results of this experiment can be found in Table 6.

This attack demonstrates that even with the phone screen off, the activities on the phone can produce consistent and detectable features in the wireless power side channel. This wireless charging attack represents a serious privacy risk as it may allow the content of a conversation or media being played to be identified even if a user does not explicitly plug in a phone to a charger.

6.4 Attacks on Cryptographic Algorithms

The power side-channel attacks on cryptographic algorithms such as RSA are widely studied and demonstrated in the context of a wired power supply. As we found that the power side channel through wireless charging is comparable to the wired power side channel for other attacks, we believe that the wireless power side-channel attacks will be feasible for cryptographic algorithms. Unfortunately, our experimental setup based on a microcontroller can only sample the power consumption every 1.4ms, which is not enough to perform a full attack on fast cryptographic algorithms. A high-end oscilloscope will be able to provide more fine-grained measurements of power consumption. As a proof-of-concept experiment, Figure 14 shows that the wireless power side channel can distinguish periods where a CPU is idle vs. repeatedly running either square or multiply operations.

7 DISCUSSIONS

7.1 Wireless vs. Wired Charging

This paper investigates the information leakage arising from the power side channel in wireless charging, using the website fingerprinting attack as a primary example. Even though a power side channel also exists in wired charging, wireless charging introduces vulnerability to attacks that would be impractical or even impossible through wired charging. In wired charging, a user needs to deliberately initiate charging by plugging a cable into a phone.

On the other hand, wireless charging can initiate without a user's deliberate actions or knowledge. If a phone is placed on a surface that contains a compatible wireless charger, then the charging process begins immediately. Wireless chargers can already be found in cars, cafes, hotels, airports, and furniture. A wireless charger could be hidden by being embedded in a surface upon which there is no symbol identifying it. In this situation, a user could place their phone on this surface, unintentionally charging it. In this sense, wireless charging can expose background computations to side-channel attacks. If the victim uses their phone without noticing it charging, foreground activity can also be exposed.

Wireless charging works over small distances or through thin surfaces. Therefore, if a phone is in a pocket or bag and a wireless charger is embedded in a seat or chair, then a charging connection could be initiated unbeknownst to an owner.

7.2 Other Use Cases of Wireless Charging Side Channel

Previous studies [10, 28] discussed how traditional power side channels may be used to detect malicious software on embedded devices. Similarly, the wireless charging interface may also be leveraged as a way to check the integrity of small mobile or embedded devices without physical connectors, such as a smartwatch. For such application scenarios, we will need further studies to see if the resolution and the accuracy of the power monitoring through wireless charging are sufficient to detect software changes or malicious activities on an embedded device.

7.3 Countermeasures

While it enables attacks without a physical connection, the wireless charging side-channel attack is still based on the same secret-dependent variations in the device's power consumption that traditional power side-channel attacks exploit. In that sense, the existing countermeasures against power side-channel attacks can also prevent the wireless charging side-channel attack. For example, Pothukuchi et al. [30] show that the power dissipated by a computer can be reshaped to obfuscate the fingerprint left by a running application. Matovu et al. [26] present both software and hardware solutions as defense mechanisms against malicious charging stations. Yan et al. [42] suggest energy obfuscation through code injection, which would embed meaningless code in applications to make features in the power trace be less predictable. Similarly, Spreitzer et al. [38] propose execution randomization as a defense mechanism against power analysis attacks. A variety of methods exist to insert random noise into a power trace or obscure sensitive information by making adjustments at the cell level [29]. Cronin

et al. [11] found that applying a low-pass filter with a cutoff of 60 Hz to collected power trace data reduced the accuracy of their passcode-cracking attack to that of a random guess.

To further reduce the amount of information leaked through wireless charging, we may be able to augment the charging algorithm to avoid fully charging the battery at less trusted locations. Previously, Zhang et al. [46] proposed a WirelessID, a system for fingerprinting individual wireless chargers and identifying potential wireless charging attacks.

Currently, iPhones running iOS 13 or later employ Optimized Battery Charging, a charging algorithm that reduces the amount of time an iPhone spends fully charged to preserve its battery lifespan. This feature uses location data to determine whether or not to delay charging past 80% [4]. If this algorithm could be adjusted to also engage when the iPhone is connected to an untrusted charger, then the battery would never leave the constant current Li-ion charging stage as seen in Figure 6. Our results show that minimal information would leak to the charger at these lower battery levels because the same amount of maximum current will be delivered to the battery regardless of the process currently executing.

8 RELATED WORK

Power analysis attacks are a well-established field of research and a variety have been studied in mobile devices. Spreitzer et al. [38] presented a thorough categorization system and survey of existing side-channel attacks, especially those applicable to mobile devices and Liu et al. [24] and Yan et al. [42] presented a survey of side-channel attacks on USB powered devices that relate to exploiting a USB connection. Clark et al. [9] found that a computer plugged into a wall was susceptible to an SPA attack and used AC power traces to carry out a website fingerprinting attack. While we build upon the existing body of power side-channel and website fingerprinting attacks to demonstrate a vulnerability, our work is the first to identify a wireless charging side channel that utilizes completely different circuitry than that of wired charging.

Genkin et al. [13] extracted ECDSA keys from a wired USB power charging side channel and also from an EM channel with a Qi charging coil as a probe. Spolaor et al. [37] showed that a malicious charging station could use a USB charging connection to exfiltrate sensitive smartphone data even when no user permissions are granted. While we also use the USB power side channel to perform attacks, we examine the wireless side channel directly via the built-in Qi-wireless charging capabilities of the phones. Additionally, we measure the current being delivered to the charger, not the physical emanations or data on the phone itself.

Yang et al. [44] determined that even when none of a smartphone's data pins are connected, a USB power station can still identify specific activity occurring on the phone. Cronin et al. [11] demonstrate that USB power traces from smartphones leak information about the contents of a device's touch screen. While we also examine this charging power side channel in our attack, our work differs in several respects. We find that the wireless side channel is as susceptible to a website fingerprinting attack as the traditional wired side channel. We also sample at 700Hz rather than 250kHz, allowing our attack to be performed by less sophisticated hardware and be more difficult to detect. Additionally, our classifier can effectively classify current traces from different device and charger models without any preprocessing.

The combination of hardware and sensor functionality on mobile devices leaves them susceptible to some unique side-channel attacks. Yan et al. [42] established a general exploitation approach for a variety of power side-channel attacks on an Android smartphone. While our attack is based on this model, we also demonstrate it on an Apple iPhone and do not require a wired connection, only the physical proximity required to wirelessly charge.

Matyunin et al. [27] successfully identify the application running on a phone by studying how CPU operations affect magnetometer measurements. Yang et al. [43] showed that the transition between running apps leaves a side channel in memory that can be used to determine what application was executing. Lifshits et al. [23] installed a malicious, power monitoring battery in a smartphone to identify various types of activity. Qin et al. [32] also adopt a similar approach to smartphone website fingerprinting by using a malicious application that estimates the fluctuation of power data. The power estimation model employs CPU data that can be accessed without permission in Android 7. In contrast to these works, our work does not require a malicious app or an otherwise compromised phone, because the act of wireless charging itself is vulnerable regardless of permissions set by the operating system.

Another method of website classification besides power side channels is through traffic and hardware analysis. In contrast to these works, our attack can occur without any software permissions at all. Hintz [16], Hayes and Danezis [15], and Lu et al. [25] measured the amount of encrypted data being transferred and other metadata to identify webpages even in the face of website fingerprinting defenses. Based on this work, Al-Shehari and Zhioua [3] proposed a unified model for traffic analysis attacks on computers. Our work also examines the Alexa top sites list [17] but differs in that the side channel exists locally on the phone's hardware, and is not a result of Internet traffic characteristics.

9 CONCLUSION

This paper presents a new side-channel attack that occurs when a Qi-compatible smart device is wirelessly charging and the power consumption of the wireless transmitter is recorded. We use a low-cost device to collect current traces and infer private information such as browser activity. We demonstrate that this attack can occur even if the user's phone is not fully charged, requires no permission from the phone OS or user, and can occur even if the acquired current trace is as short as 2.5 seconds. While this work explores a new side channel present in all wireless charging compatible smart devices, the entire scope and constraints of the wireless charging side-channel attack and useful countermeasures need to be researched in future work.

ACKNOWLEDGMENTS

We thank the Semiconductor Research Corporation (SRC) who partly supported this work through the SRC Research Scholars Program. We would also like to thank the anonymous reviewers whose comments and suggestions helped improve the clarity and quality of this paper.

REFERENCES

[1] Abien Fred Agarap. 2018. Deep Learning using Rectified Linear Units (ReLU). *arXiv e-prints*, Article arXiv:1803.08375 (March 2018), arXiv:1803.08375 pages. arXiv:1803.08375 [cs.NE]

[2] Aircharge. [n.d.]. Charger Locator App. https://www.air-charge.com/app.

[3] Taher Al-Shehari and Sami Zhioua. 2018. An empirical study of web browsers'resistance to traffic analysis and website fingerprinting attacks. *Cluster Computing* 21, 4 (2018), 1917–1931.

[4] Apple. 2020. About Optimized Battery Charging on your iPhone. https://support.apple.com/en-us/HT210512

[5] Callum Booth. 2019. Wireless charging is cool, but won't replace cables anytime soon. https://thenextweb.com/plugged/2019/01/28/wireless-charging-cables-bis-research/.

[6] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (2001), 5–32. https://doi.org/10.1023/A:1010933404324

[7] Jason Brownlee. 2018. 1D Convolutional Neural Network Models for Human Activity Recognition. https://machinelearningmastery.com/cnn-models-for-human-activity-recognition-time-series-classification/.

[8] François Chollet et al. 2015. Keras. https://keras.io.

[9] Shane S. Clark, Hossen Mustafa, Benjamin Ransford, Jacob Sorber, Kevin Fu, and Wenyuan Xu. 2013. Current Events: Identifying Webpages by Tapping the Electrical Outlet. In *Computer Security – ESORICS 2013*, Jason Crampton, Sushil Jajodia, and Keith Mayes (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 700–717.

[10] Shane S. Clark, Benjamin Ransford, Amir Rahmati, Shane Guineau, Jacob Sorber, Wenyuan Xu, and Kevin Fu. 2013. WattsUpDoc: Power Side Channels to Nonintrusively Discover Untargeted Malware on Embedded Medical Devices. In *2013 USENIX Workshop on Health Information Technologies (HealthTech 13)*. USENIX Association, Washington, D.C. https://www.usenix.org/conference/healthtech13/workshop-program/presentation/clark

[11] Patrick Cronin, Xing Gao, Chengmo Yang, and Haining Wang. 2021. Charger-Surfing: Exploiting a Power Line Side-Channel for Smartphone Information Leakage. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association. https://www.usenix.org/conference/usenixsecurity21/presentation/cronin

[12] Scott Dearborn. [n.d.]. Charging Lithium-Ion Batteries: Not All Charging Systems Are Created Equal. https://www.microchip.com/stellent/groups/designcenter_sg/documents/market_communication/en028061.pdf, last accessed on 08/28/19.

[13] Daniel Genkin, Lev Pachmanov, Itamar Pipman, Eran Tromer, and Yuval Yarom. 2016. ECDSA Key Extraction from Mobile Devices via Nonintrusive Physical Side Channels. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (Vienna, Austria) *(CCS '16)*. Association for Computing Machinery, New York, NY, USA, 1626–1638. https://doi.org/10.1145/2976749.2978353

[14] Nick Guy. 2020. Is Charging Your Phone All Day Really That Bad? https://www.nytimes.com/2020/09/04/smarter-living/phone-charging-advice.html.

[15] Jamie Hayes and George Danezis. 2016. k-fingerprinting: A Robust Scalable Website Fingerprinting Technique. In *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association, Austin, TX, 1187–1203. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/hayes

[16] Andrew Hintz. 2003. Fingerprinting Websites Using Traffic Analysis. In *Privacy Enhancing Technologies*, Roger Dingledine and Paul Syverson (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 171–178.

[17] Alexa Internet. 2020. Top Sites in United States. https://www.alexa.com/topsites/countries/US.

[18] Inviolabs. 2019. iPhone 11 Charging Test. Which USB PD Charger is better? https://www.inviolabs.com/blogs/news/iphone-11-charging-test-which-usb-pd-charger-is-better.

[19] Keith Ito and Linda Johnson. 2017. The LJ Speech Dataset. https://keithito.com/LJ-Speech-Dataset/.

[20] Serkan Kiranyaz, Onur Avci, Osama Abdeljaber, Turker Ince, Moncef Gabbouj, and Daniel J. Inman. 2019. 1D Convolutional Neural Networks and Applications: A Survey. arXiv:1905.03554 [eess.SP]

[21] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. 1999. Differential Power Analysis. In *Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO '99)*. Springer-Verlag, Berlin, Heidelberg, 388–397.

[22] John Leyden. 2018. Battery charger hack offers covert way to spy on mobile devices. https://portswigger.net/daily-swig/battery-charger-hack-offers-covert-way-to-spy-on-mobile-devices.

[23] Pavel Lifshits, Roni Forte, Yedid Hoshen, Matt Halpern, Manuel Philipose, Mohit Tiwari, and Mark Silberstein. 2018. Power to peep-all: Inference Attacks by Malicious Batteries on Mobile Devices. *Proceedings on Privacy Enhancing Technologies* 2018 (10 2018), 141–158. https://doi.org/10.1515/popets-2018-0036

[24] Hao Liu, Riccardo Spolaor, Federico Turrin, Riccardo Bonafede, and Mauro Conti. 2021. USB powered devices: A survey of side-channel threats and countermeasures. *High-Confidence Computing* 1, 1 (2021), 100007. https://doi.org/10.1016/j.hcc.2021.100007

[25] Liming Lu, Ee-Chien Chang, and Mun Choon Chan. 2010. Website Fingerprinting and Identification Using Ordered Feature Sequences. In *Computer Security – ESORICS 2010*, Dimitris Gritzalis, Bart Preneel, and Marianthi Theoharidou (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 199–214.

[26] Richard Matovu, Abdul Serwadda, Argenis V. Bilbao, and Isaac Griswold-Steiner. 2020. Defensive Charging: Mitigating Power Side-Channel Attacks on Charging Smartphones. In *Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy* (New Orleans, LA, USA) *(CODASPY '20)*. Association for Computing Machinery, New York, NY, USA, 179–190. https://doi.org/10.1145/3374664.3375732

[27] Nikolay Matyunin, Yujue Wang, Tolga Arul, Kristian Kullmann, Jakub Szefer, and Stefan Katzenbeisser. 2019. MagneticSpy: Exploiting Magnetometer in Mobile Devices for Website and Application Fingerprinting. In *Proceedings of the 18th ACM Workshop on Privacy in the Electronic Society* (London, United Kingdom) *(WPES'19)*. Association for Computing Machinery, New York, NY, USA, 135–149. https://doi.org/10.1145/3338498.3358650

[28] Jungmin Park, Fahim Rahman, Apostol Vassilev, Domenic Forte, and Mark Tehranipoor. 2019. Leveraging Side-Channel Information for Disassembly and Security. *J. Emerg. Technol. Comput. Syst.* 16, 1, Article 6 (Dec. 2019), 21 pages. https://doi.org/10.1145/3359621

[29] Thomas Popp, Stefan Mangard, and Elisabeth Oswald. 2007. Power Analysis Attacks and Countermeasures. *IEEE Design & Test of Computers* 24 (2007).

[30] Raghavendra Pradyumna Pothukuchi, Sweta Yamini Pothukuchi, Petros G. Voulgaris, and Josep Torrellas. 2019. Maya: Falsifying Power Sidechannels with Operating System Support. *CoRR* abs/1907.09440 (2019). arXiv:1907.09440 http://arxiv.org/abs/1907.09440

[31] Qi-Wireless-Charging.net. 2020. Qi Enabled Phones & Compatible Devices 2020. https://qi-wireless-charging.net/qi-enabled-phones/.

[32] Yi Qin and Chuan Yue. 2018. Website Fingerprinting by Power Estimation Based Side-Channel Attacks on Android 7. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. 1030–1039. https://doi.org/10.1109/TrustCom/BigDataSE.2018.00145

[33] Caleb Rule. 2020. Average Load Times, Core Web Vitals, & More: Fortune 100 Website Speed Tests. https://www.pedowitzgroup.com/fortune-100-websites/.

[34] Lamyaa Sadouk. 2019. CNN Approaches for Time Series Classification. In *Time Series Analysis - Data, Methods, and Applications*, Chun-Kit Ngan (Ed.). IntechOpen. https://doi.org/10.5772/intechopen.81170

[35] Tara N. Sainath, Oriol Vinyals, Andrew Senior, and Haşim Sak. 2015. Convolutional, Long Short-Term Memory, fully connected Deep Neural Networks. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 4580–4584. https://doi.org/10.1109/ICASSP.2015.7178838

[36] Laura Silver. 2019. Smartphone Ownership Is Growing Rapidly Around the World, but Not Always Equally. https://www.pewresearch.org/global/2019/02/05/smartphone-ownership-is-growing-rapidly-around-the-world-but-not-always-equally/.

[37] Riccardo Spolaor, Laila Abudahi, Veelasha Moonsamy, Mauro Conti, and Radha Poovendran. 2016. No Free Charge Theorem: a Covert Channel via USB Charging Cable on Mobile Devices. arXiv:1609.02750 [cs.CR]

[38] Raphael Spreitzer, Veelasha Moonsamy, Thomas Korak, and Stefan Mangard. 2018. Systematic Classification of Side-Channel Attacks: A Case Study for Mobile Devices. *IEEE Communications Surveys Tutorials* 20, 1 (2018), 465–488. https://doi.org/10.1109/COMST.2017.2779824

[39] The Veloxity Team. 2015. Survey Report: Cell Phone Battery Statistics 2015-2018. https://veloxity.us/2015-phone-battery-statistics/.

[40] Lance Whitney. 2017. Why You Shouldn't Charge Your Mobile Phone Overnight. https://time.com/4949569/mobile-phone-charge-overnight/.

[41] Wireless Power Consortium. 2017. *The Qi Wireless Power Transfer System Parts 1 and 2: Interface Definitions Version 1.2.3.* Available at https://www.wirelesspowerconsortium.com/knowledge-base/specifications/download-the-qi-specifications.html.

[42] Lin Yan, Yao Guo, Xiangqun Chen, and Hong Mei. 2015. A Study on Power Side Channels on Mobile Devices. In *Proceedings of the 7th Asia-Pacific Symposium on Internetware* (Wuhan, China) *(Internetware '15)*. Association for Computing Machinery, New York, NY, USA, 30–38. https://doi.org/10.1145/2875913.2875934

[43] Li Yang, Teng Wei, Jianfeng Ma, Shui Yu, and Chao Yang. 2018. Inference Attack in Android Activity based on Program Fingerprint. In *2018 IEEE Conference on Communications and Network Security (CNS)*. 1–9. https://doi.org/10.1109/CNS.2018.8433169

[44] Qing Yang, Paolo Gasti, Gang Zhou, Aydin Farajidavar, and Kiran S. Balagani. 2017. On Inferring Browsing Activity on Smartphones via USB Power Analysis Side-Channel. *IEEE Transactions on Information Forensics and Security* 12, 5 (2017), 1056–1066. https://doi.org/10.1109/TIFS.2016.2639446

[45] Gabriel Zaid, Lilian Bossuet, Amaury Habrard, and Alexandre Venelli. 2019. Methodology for Efficient CNN Architectures in Profiling Attacks. *IACR Transactions on Cryptographic Hardware and Embedded Systems* 2020, 1 (Nov. 2019), 1–36. https://doi.org/10.13154/tches.v2020.i1.1-36

[46] Jiayu Zhang, Zhiyun Wang, Xiaoyu Ji, Wenyuan Xu, Gang Qu, and Minjian Zhao. 2021. Who is Charging My Phone? Identifying Wireless Chargers via Fingerprinting. *IEEE Internet of Things Journal* 8, 4 (2021), 2992–2999. https://doi.org/10.1109/JIOT.2020.3024572

A FULL WEBSITE LIST

Table 7: Alexa top 50 sites used in Section 5.2. All connection types utilize HTTPS.

Websites	
google.com	adobe.com
youtube.com	salesforce.com
amazon.com	espn.com
yahoo.com	apple.com
facebook.com	cnn.com
zoom.us	wellsfargo.com
reddit.com	intuit.com
bing.com	nytimes.com
wikipedia.org	craigslist.org
ebay.com	slack.com
office.com	aliexpress.com
chase.com	homedepot.com
live.com	imdb.com
microsoft.com	msn.com
netflix.com	capitalone.com
instagram.com	hulu.com
zillow.com	yelp.com
twitch.tv	paypal.com
walmart.com	americanexpress.com
linkedin.com	spotify.com
force.com	usps.com
etsy.com	aws.amazon.com
dropbox.com	tiktok.com
twitter.com	alibaba.com
indeed.com	bestbuy.com