

What Do Your Smart-Home Devices Reveal About You?

Hima Boddupalli
University of Colorado Boulder
Boulder, Colorado
hima.boddupalli@colorado.edu

Shivakant Mishra
University of Colorado Boulder
Boulder, Colorado
mishras@colorado.edu

Mohammed Al-Mutawa
Kuwait University
Kuwait
almutawa@cs.ku.edu.kw

ABSTRACT

With increasing popularity of smart home devices, there is a growing concern about the security and privacy risks these devices pose to the users. Indeed, researchers have been addressing these risks over the last decade by identifying vulnerabilities and developing solutions. Despite this effort, security breaches and privacy exposures still happen routinely. While most of these privacy exposures are attributed to incorrect system configurations, this paper explores the question whether these smart home systems pose any privacy risks even when they are configured correctly and the adversary has only some very minimal capabilities. In particular, the paper presents an experimental study of two very popular smart home devices with an adversary that can passively observe all the message traffic coming in or going out of a home without being able to decrypt the messages or disrupt the traffic in any way. The key finding is that despite this minimal adversary capability, lots of fine-grained, personal information about the users and their homes can still be inferred. Further the privacy exposure becomes much more significant if the adversary has knowledge of some basic contextual information about the home residents. The paper describes these experiments and provides a detailed analysis of the data and privacy exposure.

KEYWORDS

privacy, smart-home, side-channel, security, IoT

ACM Reference Format:

Hima Boddupalli, Shivakant Mishra, and Mohammed Al-Mutawa. 2021. What Do Your Smart-Home Devices Reveal About You?. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Home automation that began as early as Year 2000 has increasingly been becoming popular with the Internet of Things (IoT) introducing this technology into our homes by rapidly applying connectivity to everyday appliances and home features. As IoT devices become a part of our daily lives, we need to take a look at the security risks and privacy concerns this smart technology introduces into our lives.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference'17, July 2017, Washington, DC, USA

© 2021 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

IoT manufacturers collect an incredible amount of data about the users and their homes with a promise to the users that these data points are used to make their smart home experience better and more personalized. For example, consider an Amazon or Google smart speaker voice assistant. It knows where you are located, what you buy, as well as your taste in music and movies. It knows when you are home, what your voice sounds like compared to, say, your roommate's — and, if you have paired it with other smart devices and what those devices are sensing. In short, it knows a lot about you. While there is some evidence that such rich personal information helps service providers to provide better, personalized services, when in the wrong hands, it's a treasure trove of personal information that can be misused for nefarious purposes such as knowing when your house is unoccupied and safe to rob, using your credit card credential to make unauthorised purchases, or getting access to the camera feeds from your home.

Indeed, users are increasingly becoming aware of these security and privacy risks. A recent survey of people from The United States, Canada, Japan, Australia, France and the United Kingdom by Consumer International and Internet Society found that about 63% of people find connected devices to be “creepy” and 75% don't trust the way their data is shared by those devices [16]. IoT device manufacturers and service providers are also responding to these concerns by putting in appropriate safety measures, such as encrypting all messages exchanged between devices and external servers or incorporating strong authentication mechanisms for device access. To some extent, users are also following various safety tips, such as using strong passwords, disconnecting devices from the Internet when not needed, installing latest security patches, etc. However, despite these safeguards from both the users and the device manufacturers and service providers, important personal information about the users and their homes does fall into wrong hands and has been exploited for nefarious purposes [4, 17].

One common technique used by the adversaries to steal such personal information despite all the safeguards is via side channel attacks, which are based on information gained from the implementation of a computer system, rather than weaknesses in the implemented algorithm itself. Timing information, power consumption, electromagnetic leaks or even sound can provide an extra source of information, which can be exploited. For example, the power output generated by pressing a particular key on the keyboard will be different from that of another key being pressed which makes it the differentiating factor without giving away what was actually pressed. By observing the frequency of the power, it can be traced to frequently occurring alphabets. Power output is just one side channel. There are many such channels for any device. If gotten hold of, they can give an inside picture to attackers which makes it a target [10][1].

In this paper, we explore the potential for a side channel attack on smart home devices, where in an adversary only has an ability to tap into and access all the Internet traffic coming in or out of a home. This is possible if the adversary has access to the home router or the router at the Internet service provider connected to the home router. We do not assume any other capability on behalf of the adversary. In particular, the adversary cannot decrypt any messages exchanged, drop, alter or inject any messages, nor can they alter the traffic flow. The question we have explored in this paper is whether the adversary with such a limited capability can still learn something about the users and their homes.

We have experimented with two very popular home devices, a pair of cameras for indoor and outdoor monitoring (Netgear Arlo Pro 2) and a Google Home Mini (Gen 2). The indoor and outdoor cameras are connected to a cloud server and their operation can be controlled by user's smartphone or tablet. The Home Mini is powered by the Google Assistant, you can ask it questions or tell it to do things using your voice. Both of these devices are connected to their respective servers on the cloud via the Internet and the messages exchanged between the device and servers are encrypted. We have experimented with a variety of indoor scenarios ranging over different number of people at home, different types of audio and movements at home, and asking questions in different styles.

Our key finding is that while these devices employ some very good security mechanisms, lots of personal information about the users and their home can still be inferred by simply monitoring and analyzing the traffic coming in and out of their homes. In particular, an adversary can infer the exact smart devices the users are using, whether the home is currently empty, or whether the current residents are behaving erratically. Further, the knowledge of any additional contextual information such as the sleeping schedules of the residents, whether children reside at home, drinking habits of the residents, or which room(s)/floor(s) the devices are located in can allow the adversary to infer much more finer grained information about the home residents. The paper describes the details of the experiments we have conducted under a range of different scenarios and the data we have collected, provides an analysis of this data, and identifies privacy leakage based on this analysis.

The rest of the paper is structured as follows. Section 2 provides a brief overview of the related work. Section 3 describes the experimental setup. Sections 4 and 5 describe the results of all the experiments involving Home Mini and Arlo Pro respectively. Section 6 describes the privacy leakage that can be inferred from these experimental results, and finally Section 7 concludes the paper.

2 RELATED WORK

In this paper, we focus on the internet-connected smart home devices, such as internet-connected appliances, lighting, sensors, door locks, security cameras, interactive smart speakers and voice services, etc. These devices rely on cloud-based integration to provide appropriate services via an API using HTTP. These devices can typically be controlled via commands from smartphones transiting via the cloud. Over the last decade, several security and privacy vulnerabilities have been reported and discussed in literature for such devices. These include privacy risks in pairing and discovery

protocols [19] and insecure communication [6], remote spying possibilities [8, 14], over-privilege [11, 12], and end-user security and privacy concerns [20].

In this paper, we focus on side channels using message traffic analysis to infer privacy information of a home. Side channel privacy attacks have been discussed quite extensively in literature, e.g. [3] and [13] demonstrate side channel attacks on anonymity networks, and [9] and [18] use traffic fingerprinting to learn about users' Internet browsing characteristics.

Message traffic analysis has been done to infer information about potentially sensitive activities in [5]. The traffic analysis performed tries to correlate device events/actions with network activity at the time of the event/action. The devices used are second generation Nest Thermostat and second generation Nest Protect Wired. It is shown that with 88% and 67% accuracy respectively, when the thermostat transitions between the Home and Auto Away mode and vice versa, based only on network traffic originating from the device. Message traffic analysis is also used in [2] to demonstrate that despite the broad adoption of transport layer encryption, smart home traffic metadata is sufficient for a passive network adversary to infer sensitive in-home activities. The specific inferences include identifying the smart home devices and whether those devices are on or off. The devices used in this research are Amazon Echo, Belkin WeMo Switch, Orvibo Smart Socket, TP-Link Smart Plug, Nest Security Camera, Amcrest Security Camera and Sense Sleep Monitor. A traffic shaper that uses anywhere between 10 KB and 40 KB worth extra bandwidth is proposed to provide a uniform traffic rate to prevent such inferences.

3 THREAT MODEL AND EVALUATION METHODOLOGY

As mentioned earlier, our goal is to explore any potential for a side channel attack on smart home devices, where an adversary only has an ability to tap into and access all the Internet traffic coming in or out of a home. In particular, the adversary is passive in nature (similar to an ISP) in that it cannot inject any new packets in the in traffic, cannot alter the contents of any packets and cannot alter the rate of packet flow. The adversary also has no access to any of the the local area network(s) (LAN) being used in the home that is being targeted. We assume that the adversary has access to a large volume of message traffic from the past that he/she can analyze using machine learning algorithms.

To gain an insight into the potential privacy leakage from smart home devices, we chose to experiment with two very popular devices, a pair of cameras for indoor and outdoor monitoring (Netgear ArloPro 2) and a Google Home Mini (Gen 2). Both of these devices need to be connected to the Internet in order for them to function properly. Our evaluation methodology is comprised of operating each device under very specific controlled scenarios and tapping the traffic from the device to the external servers and from the external servers back to the device. We then analyze this tapped traffic to identify any unique patterns that can be exploited by the adversary to infer the scenario within the home.

Figure 1 illustrates our experimental setup. We used a laptop with Intel core i5, 4GB of RAM, running Ubuntu 18.04.4 LTS, and Wireshark [7] network protocol analyzer version 2.6.10 to collect

the data. The collected data was then analyzed using pandas[15] data analysis library. For the Google Home Mini, a Wi-Fi hotspot was setup on the laptop to act as an access point that the Google Home Mini connects to. The laptop connects to the Internet through an Ethernet port. By capturing all the data coming in or going out of the wireless interface, we were able to capture all the traffic related to Google Home Mini. We needed to use a different setup to capture the traffic of the Netgear Arlo Pro cameras. The Netgear Arlo cameras come with a dedicated base station unit, the cameras connect to the base station using Wi-Fi and the base station connects to the Internet through an Ethernet port. We connected the Ethernet port of our laptop (network protocol analyzer) to a switch port that is configured to mirror all the traffic coming from the Arlo cameras base station, and setup Wireshark to capture all the traffic of the Ethernet interface.

The next step is to take files generated by wireshark and process them using pandas. We removed all the irrelevant data such as broadcasts and multicasts, and kept only the TCP and TLS traffic. The filtered data is then used to create tables that we used to generate the plots. All our experiments were repeated at least three times.

4 EVALUATION: GOOGLE HOME MINI

For all experiments with Google Home Mini, the device was triggered about 30 seconds after each controlled scenario was setup and the traffic was recorded for a duration of two minutes. We noted that Google Home Mini remains alert for about 8-10 seconds after it is triggered with "OK Google" voice command to listen for a question from the user. Our controlled scenarios for Google Mini can be divided into three different categories: baseline scenarios, standard operating scenarios and nonstandard operating scenarios.




Figure 1: Experiment setup to tap messages to or from a smart home device.

4.1 Baseline Scenarios

The *baseline scenario category* is comprised of scenarios where the device is not triggered at all and simply stays in the background. The goal here is to see if the adversary has any possibility of inferring anything about a home when the device is not being used at all. We have experimented with three different scenarios in this category: (1) when there is complete silence at home; (2) when some soft music is playing in the background; and (3) when the device is in a natural setting where there may be some background noise or conversation. Figures 2, 3 and 4 show the number of bytes transferred as a function of time for these three scenarios.




Figure 2: Google Home: complete silence at home (no trigger)




Figure 3: Google Home Mini: soft music playing in the background (no trigger)

As we can see from these figures, there is some message exchange between the Google Home Mini and the server even when the device is not triggered. We first suspected that this indicates that the Google Home Mini was eavesdropping, collecting data even when there is no trigger, especially for scenarios two and three. However, after further inspection and analysis we realized that the Google Home Mini periodically communicates with over ten different servers even when it is in complete silence. Figure 4 shows the number of bytes being sent to the individual servers while Figures 2 and 3 show the aggregate bytes. From the three figures we notice that other than the different spikes which occur at two to five minutes intervals and have values between 1 KB and 4 KB, the Google Home Mini sends less than 500 bytes/sec of data.




Figure 4: Google Home Mini: some background noise or conversation (no trigger).

Upon further analysis, we found that the spikes are a result of regular encryption key exchanges, while the other low rate traffic is mainly comprised of keep-alive packets. As such, an adversary cannot distinguish between the three scenarios within the baseline category. However, as we will see in the next subsections the traffic rate due to keep-alive is much lower than what the Google Home Mini sends when it is triggered. Therefore, if the adversary sees this type of data/plot, he/she will know for sure that no one is interacting with the Google Home Mini, and if this situation continues for a long time, he/she may conclude that no one is at home, particularly if he/she has access to some historical data to correlate with.

4.2 Standard Operating Scenarios

The *standard operating scenario category* is comprised of scenarios where the device is used properly, i.e. it is triggered correctly with “OK Google” command and an appropriate question is asked within eight seconds after the trigger. We have experimented with five different questions in this category: (1) ask a question and get an answer; (2) ask a question and get no answer; (3) say “play some trivia”; (4) say “tell a joke”; and (5) say “I am bored”. Figures 5 and 6 show the number of bytes transferred as a function of time for these five scenarios.

From both Figures 5 and 6 we notice that a tall spike (normally above 60 KBytes) is produced whenever we trigger Google Home Mini. The shape of the plot then differs depending on what the user does after the trigger. In scenarios one and two when a normal question is asked, Figure 5 shows that when we receive an answer (top), the plot becomes wider than when there is no answer (bottom). The same is true for scenarios three, four and five shown in Figure 6, where we do receive an answer for each scenario. In scenario five




Figure 5: Google Home Mini: triggered and a questions asked; with answer received (top) and with no answer received (bottom).

(I’m bored) (bottom) where we engage with the Google Home Mini we see the widest plot.

Overall, we can make three observations from these results. First, as noted earlier, this category is clearly distinguishable from the baseline category for all scenarios. Second, the adversary can fairly well distinguish between the scenarios when no answer is received (Figure 5 bottom) and when some answer is received (all other plots of Figures 5 and 6). Finally, while it is difficult to clearly identify the exact scenario when an answer is received, the width of the spikes can be used to infer the length of the interaction time, i.e. the length of an answer.




Figure 6: Google Home Mini: triggered; play some trivia (top), tell a joke (middle) and I am bored (bottom).




Figure 8: Google Home Mini: non-standard usage; triggered with “OK Google” and followed by talking gibberish for long times (10 (top), 20 (middle) or 30 (bottom) seconds).

4.3 Nonstandard Operating Scenarios

The *nonstandard operating scenarios* category is comprised of scenarios where the device is triggered correctly with “OK Google” command, but is followed by unusual or improper usage. We have experimented with five different scenarios in this category: (1) not asking a question at all; (2) asking a very long question that is not answered; (3) talking gibberish for long times (10, 20 or 30 seconds); (4) triggering the device several times without asking any question at all; and (5) asking a proper question but only after some significant pause since the trigger. Figures 7, 8, 9, and 10 show the number of bytes transferred as a function of time for these five scenarios.




Figure 7: Google Home Mini: non-standard usage; triggered with “OK Google” and followed by silence (top), asking a very long question (bottom).




Figure 9: Google Home Mini: non-standard usage; triggered with “OK Google” and followed by repeating “OK Google” very frequently for 30 sec.




Figure 10: Google Home Mini: non-standard usage; triggered with “OK Google” and followed by a question asked after a long pause; question is answered (top), question is not answered (bottom).

These nonstandard scenarios are the most revealing and most interesting ones for Google Home Mini. Since we are triggering the device, we notice the same initial spike that we saw in the standard scenarios. There is nothing special in scenarios one and two, the plot (Figure 7) gets a little bit wider when we ask a question. We do notice that in many cases the Google Home Mini captures only a part of the question if it is longer than 10-15 seconds. This is confirmed in scenario three (Figure 8) where even though we talked for 30 seconds in one of our experiments the transfer of data lasted for less than 15 seconds. Looking at this figure, we notice that if we measure the time from the peak to the end of “hump” for the three plots we see that it almost the same at just over ten seconds.

Scenario four (Figure 9) is one of the interesting ones as it is very distinct. The width of the “hump” is 30 seconds, very long compared to the previous scenarios which last on average 15 seconds. Also the values in the plot are relatively high hovering around 40KB/sec. Even for scenarios where there is a long back and forth interaction with Google Home Mini, the values drop when the user stops talking and an answer is fetched from the server. So if the

adversary notices this type of plot he/she can infer that either the user is being silly, e.g. like a child, or might be intoxicated. The fifth Scenario (Figure 10) is also revealing. At first glance, it looks similar to the standard scenarios, but upon further inspection, we notice something different. After the initial spike the values drop to about 40 KB/sec and stays there for about six seconds which is the amount of time we waited before asking the question and once we start talking/asking, it goes up by 10KB/sec.

Overall, we make three important observations from our experiments in this category. First, this category is clearly distinguishable from the baseline category for all scenarios. Second, this category can be fairly well distinguished from the standard operating scenarios category as the width of the plots are generally longer. However, there are some exceptions here, e.g. scenarios reported in Figure 7. Finally, silly behaviors like repeatedly saying “Ok Google” or talking gibberish after triggering the device are clearly distinguishable.

5 EVALUATION: NETGEAR ARLO PRO

There are two modes in this camera - armed (motion or sound detection on) and disarmed (detection off). The two modes can be set using the Arlo Pro app. A user can also live stream the video captured on his/her phone as and when needed. When the camera is armed, the user can choose between motion detection using the camera IR sensor or sound detection using the microphone. The user can choose to get a notification every time some motion or sound is detected.

As in the case of Google Home Mini, for all experiments with Arlo Pro camera, the device was triggered about 30 seconds after each controlled scenario was setup and the traffic was recorded for a duration of two minutes. Our controlled scenarios for Arlo Pro camera can be divided into four different categories: baseline scenarios, audio and motion interaction scenarios, field-of-view scenarios, and notifications scenarios.

5.1 Baseline Scenarios

The *baseline scenario category* is comprised of a scenario where there is absolutely no activity at home—no motion or sound of any kind and the app notification is off. The goal here is to see what kind of traffic is produced by the camera and if the adversary has any possibility of inferring anything about a home when there is absolutely no activity at home, a possible indication that there is no one at home. Figure 11 shows the message traffic for this scenario. As we can see in the figure, there is a periodic pattern of TCP keep-alive packets (66 Bytes in size), which are essentially used to detect when and if the camera goes offline.

5.2 Audio and Motion Scenarios

The *audio and motion scenario category* explores the impact of audio and motion on the network traffic between the camera and the servers. We have experimented with five different scenarios here: (1) there is no movement in the field of view and there is no sound; (2) there is no movement in the field of view but there is some sound (music playing in the background); (3) there is movement in the field of view and there is no sound; (4) there are movements in the field of view and there is sound as well; and (5) there is intermittent movement in the field of view and there is no sound.




Figure 11: Message traffic when the camera is not triggered

As we mentioned earlier, the camera can either be triggered by motion or by sound. In all five scenarios that we experimented with in this category, the camera was triggered using sound by tapping on the camera microphone. We chose the sound trigger to minimize the differences between the runs and between the relevant scenarios. We also used a video for the scenarios with motion to reduce the variations. Figures 12, 13, and 14 show the message traffic for these scenarios.




Figure 12: Message traffic when the camera is triggered but there is no motion, the blue line represents no sound scenario and the orange line represents the scenario with sound.

In Figure 12, we notice something very interesting; scenarios one and two have almost identical plots. In addition, after the initial large spike they have a very distinct saw-tooth shape. We conjecture that this shape is produced because of the way the video/audio stream is compressed and sent. In both scenarios, the camera is capturing the same still image that doesn’t change for the whole duration of the experiment. Also in both scenarios the camera is recording sound, in scenario one there is no sound and in scenario two music is playing in the background. Since data size to capture sound is relatively much smaller compared to pictures, the difference in the amount of data that needs to be sent between the two scenarios is so small that the plots overlap majority of the times, and because of video compression, the plots keep oscillating between 110KB/s and 40KB/s.

Figure 13 shows scenarios three and four when there is motion. We notice that again sound has no real impact on the shape of the




Figure 13: Message traffic when the camera is triggered and there is motion, the blue line represents no sound scenario and the orange line represents the scenario with sound.




Figure 14: Message traffic when the camera is triggered and there is intermittent motion

plots but since the camera in capturing motion the amount of data that need to be compressed and sent in each frame is higher than that of scenarios one and two and is mostly close to the 100KB/s level.

Figure 14 shows the plot when we start with the camera capturing video with no motion and then something moves in the field of view of the camera for a short period and then move out of the field of view. In our experiment, we introduced movement that lasted for about 10 seconds. We can clearly identify this movement in the figure by the disruption in the saw-tooth pattern near the middle of the plot.

Overall, we make three important observations. First, this category is clearly distinguishable from the baseline category, i.e. presence of motion or sound or both is clearly identifiable from when there is no sound or motion at all. Second, the presence or absence of sound cannot be identified (except the baseline category when the device not triggered). Finally, presence and duration of motion can be clearly identified.

5.3 Field of View Scenarios

The *field-of-view scenario category* is comprised of natural home settings when there are some people at home, in which case it is expected that they will come in the field of view of the camera. We have experimented with two scenarios here: (1) two people walking close to the camera across its field of view; and (2) two

people walking far away from the camera across its field of view. Figure 15 show the message traffic for these scenarios.




Figure 15: Message traffic when two people are moving in front of the camera: Top plot is when they move very close to the camera and bottom plot is when they are moving far away from the camera.

Looking at the two plots in Figure 15, we notice that when the movement is close to the camera (top plot), the saw-tooth pattern is disrupted near the beginning as oppose to the bottom plot where it is not disrupted. We reason that when the movement is close to the camera the whole video frame is changing thus more data need to sent, whereas when the movement is far very small portion of the video frame is changing requiring less data to be sent. This indicates that any motion that is far from the camera is indistinguishable from no motion at all.

5.4 Notifications Scenarios

One feature that the Arlo Pro camera, and indeed most home security cameras provide is that users can choose to receive email alerts and push notifications whenever any motion or sound is detected. The *notifications scenario category* explores if an adversary can infer whether or not the user has turned the notifications on, which would potentially indicate whether the user is at home or not. We have experimented with several identical scenarios where the only difference was that notification was on in one and off in the other. In all our experiments, we were unable to find any differences in the data/plots that would allow us to differentiate between the two scenarios of notifications on or off. We also tried to see if there is a set of servers that show up only when the notification is turned on but that was inconclusive as well. We have chosen not to include those plots in the paper for brevity.

6 PRIVACY LEAKAGE

In this section, we discuss the findings from Sections 4 and 5 and what privacy information is leaked using these or similar devices. With increasing awareness of security and privacy vulnerabilities, encryption of message content between smarthome devices and servers is now commonplace. Typically security protocols used are TLS/SSL. However, it is important to note that while the content of the messages exchanged between the devices and the servers is encrypted, their headers are not. Thus, an adversary can infer the type, make and sometimes even the model of smarthome devices that a user is using at home. This could be done either by using reverse DNS of server addresses or by simply knowing the list of IP addresses that the smart home devices connect to. For example

the Arlo Pro cameras connect to several amazon EC2 servers with some of the servers having "Arlo" in their name.

Once an adversary knows the identity of these smart home devices, they can look for usage patterns to infer the possible activity going on at home. The first privacy leakage that we discovered in our experiments is that an adversary can infer whether there is someone at home or not with fairly high probability. Note that when there is no one at home, the devices will not be triggered, no one will say "OK Google" or tap the camera, and there will be no (or minimal) sound or motion. This situation corresponds to the baseline categories that we experimented with for both Google Mini and Arlo Pro camera and found that baseline categories in each device is clearly distinguishable from all other categories. Thus by noticing a traffic pattern in line with the baseline categories, the adversary would infer that there is no one at home. We do note that it is possible that there may be someone at home but does not trigger Google Mini or come in the camera view. This would be a baseline category scenario but the adversary would incorrectly infer that there is no one at home. However, the chances are relatively low for this situation unless the people at home are sleeping, e.g. in the night. In this respect any additional clues such as day vs night or knowledge of sleeping schedules of the residents in conjunction with observations from message traffic patterns from multiple devices would further help the adversary reinforce his/her inference about whether someone is at home or not with fairly high accuracy.

Second, if the adversary infers that there is someone at home, he/she may be able to infer much finer details of the residents' current activities with fairly high probability. This is because it is possible to infer whether the devices are being used in a normal standard way or non-standard way by observing the traffic patterns. In both the devices, we observed that non-standard usage results in a traffic pattern that is distinguishable from standard usage or the baseline categories. A non-standard usage typically indicates erratic behavior that could be attributed to either children being at home (possibly alone) or people being intoxicated. Again, any additional clues such as knowledge of whether children reside at home or residents' drinking habits in conjunction with observations from message traffic patterns from multiple devices would further help the adversary reinforce his/her inferences about the erratic behavior of the residents with fairly high accuracy.

Ability to detect non-standard usage is especially problematic as it allows the adversary to gain finer grained information about the residents' activities. We experimented with only a few non-standard usage in this paper. What other activities about the residents can be inferred from other kinds of non-standard usage is a future area of research for us.

Finally, knowledge of additional contextual clues can provide more ammunition to the adversary to infer finer grained details about the home residents. One such contextual clue is if the adversary knows the location (which room or floor) these devices are installed in. For our two smart home devices, if the adversary knows that the devices are placed in two separate rooms/floors, simultaneous triggering of both devices would indicate that there are at least two people at home. Also, if the adversary knows that a single person lives at home, he/she can infer which room/floor the resident is in

and in fact may be able to track his/her movement by observing the temporal sequence of device activations.

We note that users may switch on/off some of the devices when they are at home or when they are away depending on the nature of the device. For example, in case of indoor security cameras, users may arm/activate the camera only when they are away. Similarly, users may switch off or disconnect voice service devices such as Google Home Mini when they are away. This would certainly change the traffic patterns and possibly limit the extent of privacy leakage.

7 CONCLUSIONS

This paper explores the extent of privacy leakage from smart home devices in situations where an adversary has minimal capabilities, simply an ability to tap into and access/analyze all the Internet traffic coming in or out of a home, e.g. by having access to the home router or the router at the Internet service provider. The adversary cannot decrypt any messages exchanged, drop, alter or inject any messages, nor can they alter the traffic flow. The key finding is that even such a weak adversary can infer quite a lot of privacy information about the home and residents. With a high probability, the adversary can infer if no one is at home at present and whether the residents currently at home are behaving erratically. Knowledge of any additional contextual information such as the sleeping schedules of the residents, whether children reside at home, drinking habits of the residents, or which room(s)/floor(s) the devices are located in can allow the adversary to infer much more finer grained information about the home residents.

There are several future research directions that we plan to pursue. First, our current analysis has focused on observing the traffic pattern over time plotted as a graph. With large amount of traffic pattern data that we have been collecting, the next step would be to explore machine learning algorithms to develop classifiers that can predict with fairly high accuracy any information about home or resident activities. Second, we plan to explore additional non-standard usage of smart home devices to see if they can reveal any finer grained privacy information. Third, we plan to incorporate more smart home devices in our study and explore the extent of privacy leakage with the plethora of devices that are now becoming commonplace in our homes. Finally, while this paper is focused on detecting privacy leakage from smart home devices, the next step would be to explore possible solutions to prevent such leakage. For example, traffic shaping by sending extraneous data has been explored in literature to obfuscate network traffic patterns, but that comes at a cost of additional network bandwidth consumption. We plan to investigate other possible solutions that would involve less overhead.

ACKNOWLEDGMENTS

This research was partially supported by a grant from The National Science Foundation (CNS-1816379).

Vivek Panchal helped record some of the experiments reported in this paper.

REFERENCES

- [1] M. A. N. Abrishamchi, A. H. Abdullah, A. David Cheok, and K. S. Bielawski. 2017. Side channel attacks on smart home systems: A short overview. In *IECON 2017 - 43rd Annual Conference of the IEEE Industrial Electronics Society*. 8144–8149.
- [2] Noah Apthorpe, Dillon Reisman, Srikanth Sundaresan, Arvind Narayanan, and Nick Feamster. 2017. Spying on the Smart Home: Privacy Attacks and Defenses on Encrypted IoT Traffic. *CoRR* abs/1708.05044 (2017). arXiv:1708.05044 <http://arxiv.org/abs/1708.05044>
- [3] Adam Back, Ulf Möller, and Anton Stiglic. 2001. Traffic analysis attacks and trade-offs in anonymity providing systems. In *International Workshop on Information Hiding*. Springer, 245–257.
- [4] Confirmed: 2 Billion Records Exposed In Massive Smart Home Device Breach. <https://www.forbes.com/sites/daveywinder/2019/07/02/confirmed-2-billion-records-exposed-in-massive-smart-home-device-breach/6fd46a43411c>. accessed: 2020-07-28.
- [5] B. Copos, K. Levitt, M. Bishop, and J. Rowe. 2016. Is Anybody Home? Inferring Activity From Smart Home Network Traffic. In *2016 IEEE Security and Privacy Workshops (SPW)*. 245–251.
- [6] Ang Cui and Salvatore J Stolfo. 2010. A quantitative analysis of the insecurity of embedded network devices: results of a wide-area scan. In *Proceedings of the 26th Annual Computer Security Applications Conference*. 97–106.
- [7] Wireshark Go Deep. <https://www.wireshark.org>. accessed: 2020-07-21.
- [8] Tamara Denning, Tadayoshi Kohno, and Henry M Levy. 2013. Computer security and the modern home. *Commun. ACM* 56, 1 (2013), 94–103.
- [9] Edward W Felten and Michael A Schneider. 2000. Timing attacks on web privacy. In *Proceedings of the 7th ACM conference on Computer and communications security*. 25–32.
- [10] M. Frustaci, P. Pace, G. Aloï, and G. Fortino. 2018. Evaluating Critical Security Issues of the IoT World: Present and Future Challenges. *IEEE Internet of Things Journal* 5, 4 (2018), 2483–2495.
- [11] Grant Ho, Derek Leung, Pratyush Mishra, Ashkan Hosseini, Dawn Song, and David Wagner. 2016. Smart locks: Lessons for securing commodity internet of things devices. In *Proceedings of the 11th ACM on Asia conference on computer and communications security*. 461–472.
- [12] Philipp Morgner, Stephan Mattejat, and Zinaida Benenson. 2016. All your bulbs are belong to us: Investigating the current state of security in connected lighting systems. *arXiv preprint arXiv:1608.03732* (2016).
- [13] Steven J Murdoch and George Danezis. 2005. Low-cost traffic analysis of Tor. In *2005 IEEE Symposium on Security and Privacy (S&P'05)*. IEEE, 183–195.
- [14] Temitope Oluwafemi, Tadayoshi Kohno, Sidhant Gupta, and Shwetak Patel. 2013. Experimental security analyses of non-networked compact fluorescent lamps: A case study of home automation security. In *{LASER} 2013 ({LASER} 2013)*. 13–24.
- [15] pandas Python data Analysis Library. <https://pandas.pydata.org>. accessed: 2020-07-21.
- [16] People say they care about privacy but they continue to buy devices that can spy on them. <https://www.vox.com/recode/2019/5/13/18547235/trust-smart-devices-privacy-security>. accessed: 2020-07-28.
- [17] The Botnet That Broke the Internet Isn't Going Away. <https://www.wired.com/2016/12/botnet-broke-internet-isnt-going-away/>. accessed: 2020-07-28.
- [18] Tao Wang, Xiang Cai, Rishab Nithyanand, Rob Johnson, and Ian Goldberg. 2014. Effective attacks and provable defenses for website fingerprinting. In *23rd {USENIX} Security Symposium ({USENIX} Security 14)*. 143–157.
- [19] David J Wu, Ankur Taly, Asim Shankar, and Dan Boneh. 2016. Privacy, discovery, and authentication for the internet of things. In *European Symposium on Research in Computer Security*. Springer, 301–319.
- [20] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End user security and privacy concerns with smart homes. In *Thirteenth Symposium on Usable Privacy and Security ({SOUPS} 2017)*. 65–80.