A False Sense of Home Security – Exposing the Vulnerability in *Away* Mode of Smart Plugs

Austin Wang and Shahriar Nirjon
Undergraduate Mobile Computing Systems Research Group
University of North Carolina at Chapel Hill
austinyw@live.unc.edu, nirjon@cs.unc.edu

***Abstract*—With smart and connected IoT systems becoming increasingly prevalent in homes and businesses, ensuring that their 'smartness' is up to the mark in making our lives secure is vital. This paper reveals the inadequacy of the 'away mode' of smart plugs–which is a feature meant to safeguard a user's home security by mimicking the on-off pattern of appliance usage when a user is away, e.g., on a vacation. Using statistical methods, we show that a third party attacker can easily extrapolate sensitive usage information by analyzing accessible data gained from these commercial, off-the-shelf smart plugs. We recommend that the developers of these smart plugs should focus on implementing rule-based, coordinated, and/or learning-based models to enhance the away-mode behavior of these smart plugs to improve their performance in obfuscating user absence.**

***Index Terms*—component, formatting, style, styling, insert**

I. Introduction

Appliance use in a home contains sensitive information that users may want to hide from the public. For example, the on/off pattern of lights indicates when the residents are at home and when they are away. Studies show that observation of appliance activity reveals user habits [2]. Outside parties or attackers are able to acquire this data by observing visual, audio, or RF indicators of appliance usage. Furthermore, with the increasing prevalence of IoT devices, wireless snooping has become a major threat. IoT products that send and receive data over the Wi-Fi are demonstrably vulnerable to data leakage via packet sniffing [4] [5].

Smart plugs are Internet-connected electric plugs that can be turned on or off remotely by a user via his or her smartphone. Presently, several commercial, off-the-shelf smart plugs provide an *'away'* mode feature that allows users to specify a time period during when their smart device will turn on and off to simulate human usage. Unfortunately, the away mode in commercial smart plugs is insufficient in hiding the habits of its users.

In this paper, we study the away mode of the two most popular commercial smart plugs to understand the algorithm behind their away mode behavior. We collect empirical data and compare human usage patterns of appliances and patterns generated by the away mode of these smart plugs. We find that the usage patterns of away modes are significantly statistically distinguishable from those of human usage.

Based on the outcomes of the study, we propose several tiers of solutions, such as creating more sensible rules for independent devices to more accurately simulate human activity, creating rules and coordination between multiple standalone devices, and enhancing these rules by learning from user activity. We leave the implementation of these solutions as our future work.

II. Away Mode of Smart Plugs

A smart plug is a plug that is connected to the Internet, and hence, it can be turned on or off remotely by a user using a smartphone.

The *away* mode is a feature advertised by many commercial, off-the-shelf smart plugs such as TP-Link Kasa [5] and Belkin Wemo [4] smart plugs. When turned on, this feature supposedly turns the device on and off randomly to simulate human activity. The advertised purpose of this feature is to make it seem as though the user is at home when they are not, thereby deterring people from vandalizing, breaking into, or stealing property. In general, away mode is a means of obscuring the usage habits of the device owner when they are away.

For example, a person who works from 9AM to 5PM on weekdays may turn on away mode during this period of absence. Another person might use away mode for an entire week while they are on a vacation in another state or country. In these cases, an attacker who does not know the homeowner's true location will not be able to tell when he is away, and will be unsure if it is safe for the attacker to commit crime.

A cost of enabling the away mode is the increased energy usage. Instead of leaving appliances off when users are away, the away mode will turn them on and off. This increases the up-time of the devices and the amount of electricity used in the house. A practical away mode should seek to limit the amount of energy consumed while making sure that it is smart enough to fool an attacker into thinking that the user is at home.

III. Vulnerability Study

We conduct an empirical study to understand the away mode behavior of commercial, off-the-shelf smart plugs. The goal of the study is to understand the underlying algorithms these devices use to simulate human activities, and we do so by observing the state changes produced by these plugs. We seek an answer to the following research question: *'is it possible for an attacker to infer that a smart plug is operating in its*

978-1-5386-9151-9/19/$31.00 ©2019 IEEE

away mode, by observing its on/off pattern over a reasonable, limited period in time?'

A. Devices

The two smart plugs we use in this study are the TP-Link Kasa HS100 [4] and the Belkin Wemo Mini [5]. We pick these two plugs for study as they are the two most popular vendors of smart devices (according to Amazon's popularity rating) who support the away mode. These smart plugs are stand-alone IoT devices that do not require a central hub. Both advertise an away mode that supposedly turns them on and off 'randomly' in order to make it seem like the user is at home when they are actually not. Belkin states that their plugs (Wemo) will remain on for a minimum of 30 minutes before turning off when they are in the away mode.

On these two devices, the start and the end time of away mode are set by the user. Additionally, these two times are required to be within the same day (between 12:00AM and 11:59PM). This suggests that the away mode in these plugs is advertised for at most a day of operation, as opposed to longer duration operations e.g., when a user is on a multiple day vacation.

B. Datasets

Since our end goal is to understand the capability of smart plugs (away mode) in imitating appliance usage by humans, we collect empirical data for both (1) a smart plug's (away mode) on/off pattern, as well as, (2) a human user's regular usage pattern of an appliance.

We consider two kinds of appliances in our study: two lights which are frequently turned on and off, and a microwave which is used occasionally. One light is setup in a bedroom and the other one in the kitchen near a microwave. The setups are shown in Figure 1 and Figure 2, respectively.

Two smart plugs are used to automatically log both the human activity and the away mode data. Additionally, human usage information from the UMass Trace Repository is used, specifically ”Electrical data from dimmable and non-dimmable switches” from Home A of the *UMass Smart* Dataset - 2013 release* [1].

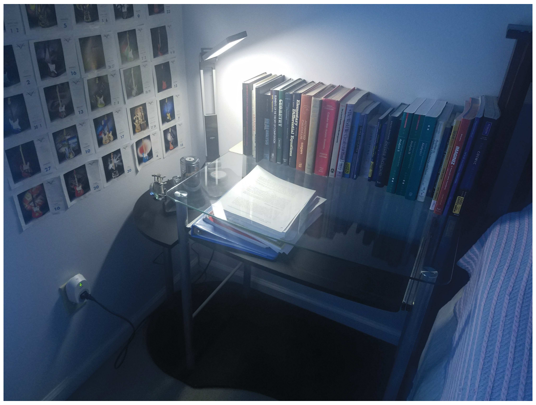

Fig. 1. The bedroom light setup

Fig. 2. The kitchen light and microwave setup

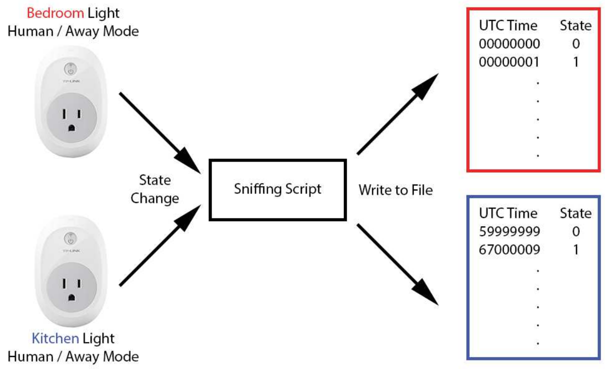

UTC Time	State
00000000	0
00000001	1
.	.

UTC Time	State
59999999	0
67000009	1
.	.

Fig. 3. A script sniffed Wi-Fi packets from each smart plug and recorded timestamps and states whenever the plug changed states

C. Data Collection

We develop two scripts to sniff Wi-Fi packets sent from the plugs. The setup is shown in Figure 3. For the Wemo plug, we use Ouimeaux API (for python) [6] to collect human and away mode activity. For the HS100 plug, we develop a node.js script using the hs100-api npm module [7] to collect human and away mode activity. Both scripts poll the plug's on/off status at one second intervals and record the UTC timestamps of state changes and on/off states at those times. The amount of time between each state change was calculated from this data.

Both smart plugs are set to away mode for different times of day to represent different use cases. For example, all day (12:00AM to 11:59PM), work day (9:00AM to 5:00PM), and evening (5:00PM to 9:00PM). All away mode times are set in the local time (Eastern Standard Time).

D. Data Analysis Tools

Additional tools are developed to analyze the data. A python script calculates the time between each off-to-on state transition for both human activity and smart plug away mode data traces. We refer to this statistic as the *inter-on time* in this paper. The script generates empirical CDFs and kernel density estimations of the inter-on time. Another python script runs Kolmogorov-Smirnov tests [9] on these inter-on time distributions to analyze their similarity.

We compute another static from pairs of appliances. For one appliance, we consider the timestamp of each off-to-on state transition, and then find the nearest off-to-on timestamp for the second appliance. Next, we take the difference of these two timestamps. We call this difference the *closest inter-on time* between two appliances. We generate the empirical CDFs and kernel density estimations from the closest inter-on times.

IV. Observations from the Study

In this section, we describe our observations from the study.

A. Predictability of Away Mode

An attacker may find that the on/off pattern of an appliance usage is predictable and thus determine that it is not a human activity. The plots of on/off states of Wemo and HS100 smart plugs are shown in Figure 4 and Figure 5, respectively.

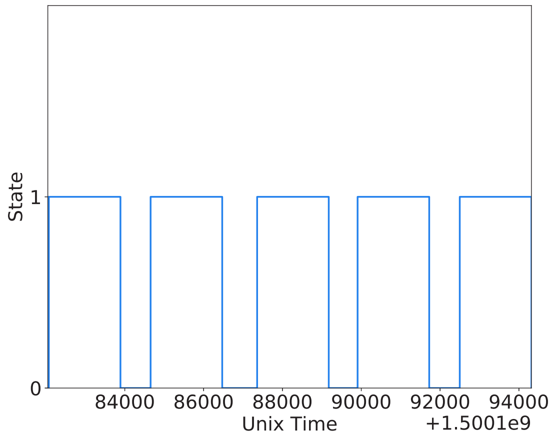

Fig. 4. State plot of the Belkin Wemo Mini's away mode shows a predictable pattern between state changes.

The Wemo Mini's away mode display a highly predictable pattern in its times between state changes. The plug powers on for periods of exactly 30m:17s or 30m:18s before turning off for up to 15 minutes. This pattern repeats until the away mode ends. This pattern is very predictable – meaning an attacker would easily identify when the Wemo Mini is in away mode by observing its constant up time.

The HS100 does not present an immediately predictable pattern in its time between state changes. However, Wi-Fi snooping on the HS100's packets reveal that the device sends information about its away mode settings including whether or not it is active. While this is by itself a major security flaw, we continue experimentation on the HS100. We continue to examine the predictability of the HS100 in the case that an attacker only has snooping capabilities (visual and acoustic cues) to detect state changes. Another observation is that the HS100 changes its state exactly 6 times over the duration of away mode, regardless of the start and end times. This means that shorter periods of away mode result in more frequent state changes. Observing these patterns alone, however, does not provide enough data to differentiate between human activity and HS100's away mode.

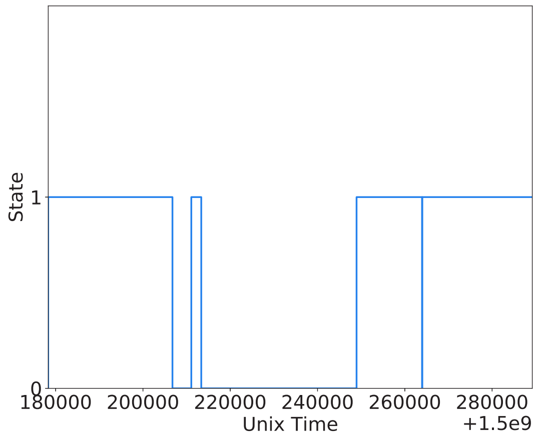

Fig. 5. State plot of the TP-Link Kasa HS100's away mode.

B. Individual Appliance Activity

An attacker may perform statistical analysis on a single appliance usage to determine whether it is human operated or it is a smart plug acting in its away mode.

When comparing inter-on times from two separate weeks of human data with a K-S test, we are unable to say that the two sets of data come from different distributions. In other words, the analysis suggests that a person's appliance usage should be statistically highly similar from week to week.

However, the same analysis between human data and the Wemo Mini's away mode data show a statistically significant difference. This means that away mode displays usage patterns that are unlikely to be generated by human activity. An attacker would be able to detect when away mode is on simply by collecting a single appliance's data and running statistical analysis.

Samples Tested	KS test statistic	p-value
Human vs Human	0.20779	0.80651
Human vs Wemo	0.95455	1.58812e-13

TABLE I
Individual Plug Activity

For instance, Figure 6 and Figure 7 show the probability density functions of human generated and a smart plug generated inter-on times, respectively. We observe that the distributions are very different. Table I provides the KS statistics for the two cases: human vs. human and human vs. Wemo. A higher KS score and a lower p-value mean that the distributions are different and the results are reliable (high confidence). Hence, an attacker who observes the usage pattern of a target appliance and has the usage pattern of a regular

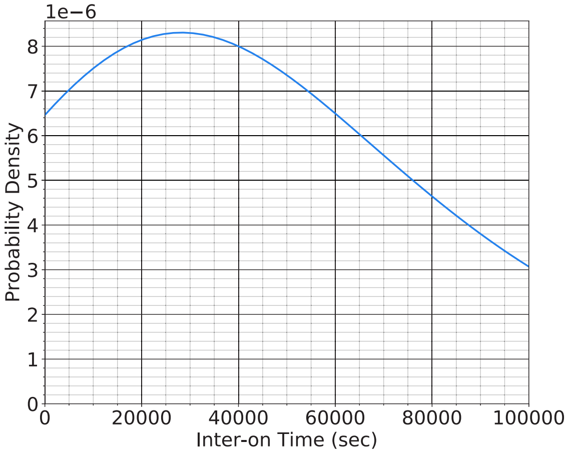

Fig. 6. For individual appliance: Probability distribution produced by Gaussian KDE of human generated inter-on times.

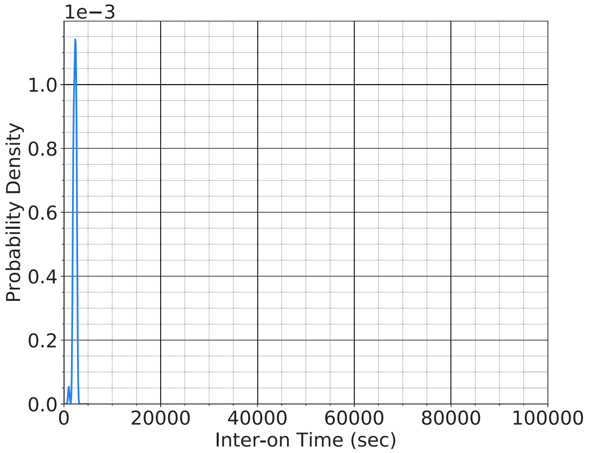

Fig. 7. For individual appliance: Probability distribution produced by gaussian KDE of the Belkin Wemo Mini's away mode inter-on times.

human being from an external source can use KS statistic to determine if the target appliance is used by a human or it is controlled by the away mode.

C. Coordinated Appliance Activity

An attacker may perform statistical analysis on the relationship between usages of two appliances to determine whether they are human operated or smart plugs acting in their away mode.

Taking the closest inter-on time of two appliances results in a distribution that describes a relationship between these appliances. In this analysis, we use human data from the UMass Trace Repository. We take two separate weeks of human generated data and compute the closest inter-on times between a bedroom light and a kitchen light. After generating the distributions using ECDF and Gaussian KDE, we compare the two weeks and find no significant evidence that they come from different distributions. Again, this indicates that the coordination/rule between these two appliances occur from week to week in human activity. In other words, there is a pattern in using two appliances by a human user, which is consistent over multiple weeks.

The analysis is performed again with the same appliances but this time we compare the closest inter-on time of human usage with the HS100's away mode's. We find that there is a statistically significant chance that the two samples come from different distributions. This suggests that the HS100 does not follow the same rule-based patterns that applies to human appliance usage.

This deviation suggests a flaw in the away mode of the HS100. An attacker can detect when a plug is in away mode by collecting state data of two coordinated appliances and then determine its operating mode with a high degree of confidence through statistical analysis.

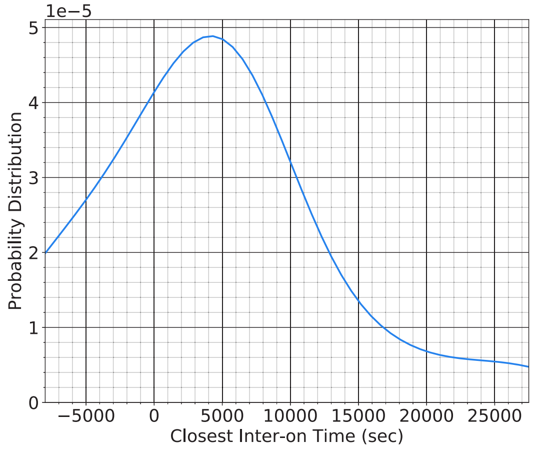

Fig. 8. Multiple appliances: The distribution generated by the Gaussian KDE of the closest inter-on times between human usage of bedroom and kitchen lights.

Samples Tested	KS test statistic	p-value
Human vs Human	0.28947	0.19899
Human vs HS100	0.63158	0.00783

TABLE II
COORDINATED PLUG ACTIVITY

Figure 8 and Figure 9 show the distributions of the closest inter-on times between pairs of appliances when they are used by humans and generated by the away mode, respectively. Table II shows the KS test results along with the p-values. We observe that the distributions are very different and by performing the KS test on the distributions of closest inter-on times, an attacker can determine with a high confidence if the usage pattern reflects that of a human or a smart plug (away mode).

V. INTELLIGENT AWAY MODE DESIGN

Our study suggests that the away modes of smart plugs are insufficient in protecting a user's privacy. The Wemo's

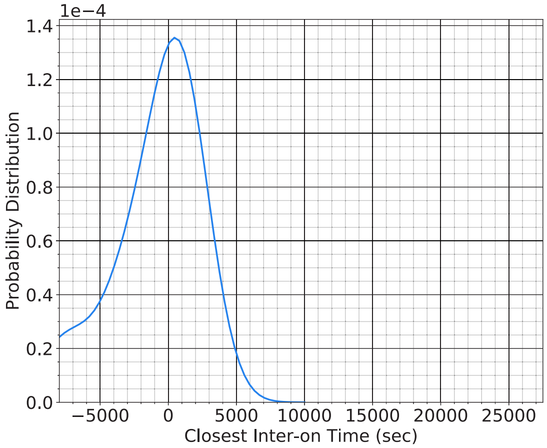

Fig. 9. Multiple appliances: The distribution of the HS100's closest inter-on times shows a long left tail.

away mode shows a lack of randomness, making it simple for an attacker to detect through observation alone. Furthermore, its patterns are not sufficiently similar to human activity, so comparing any one appliance could reveal the difference between away mode and human usage. The HS100 reveals its away mode schedule through insecure packets, making its away mode extremely vulnerable to Wi-Fi sniffing. If an attacker only has access to visual/audio/RF cues to detect state changes, he can determine when the HS100 is in away mode by watching two appliances that are used in some coordinated manner. Based on these observations, we recommend the following criteria to consider when designing an intelligent away mode for the smart plugs:

- Firstly, *the predictability of individual plugs needs to be decreased*. Patterns such as those observed in the Wemo Mini's away mode are far to unnatural. The randomization of state changes needs to be good enough that an attacker cannot easily detect the away mode just by examining the state changes. A few things may help achieve this goal. Smart plugs should consider the duration during which they are active and adjust accordingly. For example, while the user may turn the lights on many times over the course of a day, they would not turn them on as often over the course of an hour. Additionally, giving smart plugs some basic pre-calculated model of human activity can allow away mode to better simulate activity even with original manufacturer settings.
- Secondly, *away mode effectiveness can be increased by coordinating plug activity*. Some pairs of appliances may be turned on together frequently while other pairs may rarely be used together. In order to improve the accuracy of simulations, smart plugs should be able to communicate certain information to each other. Future work can explore energy efficient ways to establish hub less communication between smart plugs.
- Thirdly, *away mode can better obscure user habits if it can be personalized to each user*. Each person's appliance usage is different and it may even evolve over time. Further research should include machine learning to adapt away mode patterns to better fit specific users and changing habits.

VI. Possible Solutions

No current research meets the necessary requirements for a secure and cost-effective implementation of away mode. Nevertheless, they offer valuable tools and approaches to develop better models of simulating human activity.

A. Agent-Based Simulation

One method of designing an intelligent away mode is Brahms agent-based simulation of human activity [8]. This model considers actions, such as appliance usage, to be activities enacted by an agent. These activities are split between deliberate activities that are taken by the agent to achieve some goal and reactive activities that occur in response to the environment. Additionally, these activities may depend on each other. For example, some activities are more likely to occur together or some activities may interrupt others. When simulating multiple users, this method requires some sort of communication between agents in order to model collaborative activities, such as coordinated away mode.

While Brahms agents can be effective, there are limitations in its application to away mode for IoT devices. Firstly, it requires many initial parameters such as facts about the environment and information about the user. Obtaining information and catering it to specific users is not feasible in commercial devices that are used by many people and in many different settings. Furthermore, simulating agents requires constant computation which may greatly increase the energy consumption of smart plugs. Energy consumption is the major cost of away mode and should be kept as low as possible. Finally, continuous communication is important in agent-based modelling, but hub-less smart plugs have limited communication capabilities.

B. Situation Models

Another way of simulating activity is through situation models [3]. Like agent-based simulation, this method uses the environment and context clues to inform the model. The naive approach would be to use a model with fixed context. This would clearly not work because a variety of people with different habits need to be able to use away mode. The proposed approach is to use machine learning to adapt to new and evolving environments. Using 3D video tracking system and ambient sound and speech detection, researchers trained the model to recognize object roles and relations between objects.

The biggest drawback of situation models is the requirement for sensors. In order to effectively contextualize the environment, cameras and sound sensors need to be positioned to capture 3D position and other relevant information. This introduces new costs to smart device away mode. Not only would the user be required to purchase extra hardware sensors with their smart plugs, they would also have to install these sensors in rooms where they want to use away mode. Additionally, this introduces much higher electricity costs because these sensors need to be on frequently to collect data. Furthermore, new sensors may also introduce new points of failure in the security of an IoT home. Video/audio sensors can provide especially sensitive information. If an outside party were to gain access to video footage of a room, they would also be able to determine user habits, thereby rendering away mode useless.

VII. Conclusion

If smart and connected homes are to become prevalent, we must ensure that the privacy of end users is protected. Fixing security flaws is a priority, but we should also seek to develop features that enhance privacy. Away mode of IoT devices is one such promising feature. If implemented correctly it can obfuscate user habits from outside parties. The key goal is to cause uncertainty about whether the state transitions are produced by human usage or away mode. If this condition is met, it will be difficult for an attacker to determine user habits.

References

[1] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht. Smart*: An open data set and tools for enabling research in sustainable homes. SustKDD, August, 111:112, 2012.

[2] K. Basu, L. Hawarah, N. Arghira, H. Joumaa, and S. Ploix. A Prediction System for Home Appliance Usage. Energy and Buildings, 67 (Supplement C):668 679, 2013.

[3] O. Brdiczka, J. L. Crowley, and P. Reignier. Learning Situation Models in a Smart Home. IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics, 39(1):5663, Sept. 2008.

[4] G. Georgovassilis. Controlling the TP-Link HS100 Wi-Fi Smart Plug, January 2017.

[5] P. Joaquin. Belkin Wemo Switch NMAP Scripts, June 2017.

[6] I. McCracken. GitHub Project Page of Ouimeaux. https://github.com/iancmcc/ouimeaux, October 2018.

[7] P. Seal. GitHub Page for HS100 API. http://github.com/plasticrake/hs100-a, 2017.

[8] M. Sierhuis, W. J. Clancey, R. van Hoof, and R. de Hoog. Modeling and Simulating Human Activity. May 2000.

[9] R. Wilcox. Kolmogorov-Smirnov Test. John Wiley Sons, Ltd., 2005.