Modelling Traffic Analysis in Home Automation Systems

Frederik Möllers[1], Stephanie Vogelgesang[2], Jochen Krüger[1], Isao Echizen[3], and Christoph Sorge[1]

[1] CISPA, Saarland Informatics Campus
[2] Ministry of Justice, Saarland
[3] National Institute of Informatics, Tōkyō

Abstract The threat of attacks on Home Automation Systems (HASs) is increasing. Research has shown that passive adversaries can detect user habits and interactions. Despite encryption and other measures becoming a standard, traffic analysis remains an unsolved problem. In this paper, we show that existing solutions from different research areas cannot be applied to this scenario. We establish a model for traffic analysis in Home Automation Systems which allows the analysis and comparison of attacks and countermeasures. We also take a look at legal aspects, highlighting problem areas and recent developments.

1 Introduction

HASs are an emerging trend in consumer electronics. The benefits are promising: an increased comfort of living; savings on energy and resource consumption; increased safety and security. However, HASs have been developed with a focus on usability, energy efficiency and low cost. In addition to active attacks, adversaries can passively intercept communication using cheap, available hardware. Smart homes can thus actually facilitate privacy breaches.

Encryption and other methods do not completely solve this problem. Traffic analysis attacks disclose habits as well as presence or absence of users. In order to counter this, dummy traffic can be generated by the system. However, generating too much dummy traffic negatively affects the lifetime of battery-powered devices and can exceed regulatory thresholds.

Our contributions in this paper are as follows:

- We establish system and attacker models for traffic analysis attacks in Home Automation Systems. They allow modelling realistic attack scenarios such as those shown in previous works and are extensible to account for new findings.
- We formulate privacy goals for Home Automation Systems using ideas from the field of Private Information Retrieval. By building on established definitions, we can leverage research that has been conducted in related fields.
- We illustrate the application of our definitions of privacy goals by examining two approaches to dummy traffic.

The final publication is available at Springer via
https://doi.org/10.1007/978-3-030-02641-7_27.

– We sketch issues and current developments in the interaction between technology and legal frameworks (especially criminal law). We briefly describe how technology and international law are intertwined and can work together as a comprehensive concept for data protection.

2 Related Work

Research so far has tackled related problems in different scenarios. For HASs, problems have been identified but no solutions have been proposed so far.

2.1 Wireless Sensor Networks

In both Wireless Sensor Networks (WSNs) and HASs, devices have little computational power and communication is costly in terms of power consumption. However, research on privacy in WSNs mainly focuses on location privacy [1,2,3] instead of hiding the existence of communication.

Yang et. al. have proposed a scheme employing constant-rate dummy traffic generation which could be applied to systems not using multi-hop routing. [4] We examine such a scheme in Sec. 5.2. The authors also propose a scheme for random dummy traffic generation requiring much less overall traffic [5], but this introduces delays which are to be avoided in HASs. Furthermore, the behaviour of the inhabitants might not conform to the distribution assumed in their work.

2.2 MIX networks

Early approaches of dummy traffic generation [6] use constant-rate traffic, which we examine in Sec. 5.2. More recently, there has been significant research on traffic analysis in low-latency MIX networks. [7] Shmatikov and Wang [8] have developed an approach which aims to find a balance between the amount of dummy traffic and the success rate of traffic analysis attacks. Their approach may be applicable to HAS networks with modifications, though their evaluation is based on HTTP traffic samples. It is unclear how much of the approach can be applied to HASs and whether the same goals can be achieved.

2.3 Differential Privacy

Definitions of Differential Privacy are used to develop techniques which provide unobservability of events or user data. However, ε-Differential Privacy is a property of a specific function [9]. In our scenario we do not know the informational value of specific data points and we do not know which computations an attacker will perform on captured traffic.

Dwork et. al. have developed an approach to continuously monitor an event source and count its events with the counter guaranteeing ε-Differential Privacy even if its internal state is visible to the attacker at some point in time. [10] However, the guarantees do not necessarily apply to other functions.

2.4 Steganography and Covert Channels

Steganography can hide communication without generating any dummy traffic. However, since there is no cover data available in HAS communication, many steganography approaches cannot be applied to this setting. Hiding the traffic among noise of a wireless channel has been investigated by Bash et. al. [11] It is unclear whether approaches on higher layers of the network stack can either exceed some fundamental limitations or achieve the same results at lower (manufacturing or communication) costs. Furthermore, the approach was developed for wireless networks and might not be applicable to wired systems.

2.5 Home Automation Systems

Möllers et. al. have shown that unencrypted, wireless HAS communication leaks automation rules and user habits to passive observers. [12] Mundt et. al. have shown that wired systems are also susceptible to the same eavesdropping attacks. [13] Encryption and padding do not protect against statistical disclosure attacks, as was shown by Möllers et. al. [14] Other authors have focused on other aspects of security, for example confidentiality of information and access controls. [15]

3 System and Attacker Model

In this section, we first list our assumptions and then describe the model.

3.1 Assumptions

System: Network Topology and Routing We assume that the network graph is a clique with respect to intended communication, i.e. no routing is necessary. The reason behind this assumption is that routing introduces a set of problems, but also opportunities which are already well understood.

System: Encryption We assume that the HAS uses padding and encryption for both message payloads and addressing information. The attacker cannot break the encryption in reasonable time and is unable to learn information about the sender, receiver or contents of a message. If state-of-the-art approaches are applied correctly, this assumption is reasonable. Even though we do not know of any system actually using both full packet encryption and padding, we consider this to be an engineering problem rather than a research one.

Attacker: Mode of Operation The attacker in our scenario is passive. No active attacks such as traffic injection, node compromise or DoS are launched.

Traffic injection does not lead to a reaction by the system if encryption and authentication are correctly implemented. Node compromise requires either the presence of vulnerabilities in the nodes' software or physical intervention by the attacker. Denial of service attacks offer no benefit for the attacker (messages are resent after the attack ends) and are detectable.

Attacker: Reception The attacker has perfect reception as well as an accurate clock. They are able to capture all traffic, do not suffer from reception errors and can save the time at which a message was captured. Both experiments on real-world installations of HASs [12,13] have proven this to be realistic.

Attacker: Limits The attacker does not launch triangulation or device fingerprinting attacks. [16] These attacks require a considerable amount of effort from the attacker and countermeasures are a separate area of research.

Attacker: Awareness and Knowledge The attacker is aware of the underlying algorithms of countermeasures, but does not know runtime information (e.g. the internal state of PRNGs). We model privacy goals with respect to a given set of tasks that the user might perform. These may differ vastly depending on the setting, so any assumption reduces the utility of the model.

3.2 System Model

Communication packets from the HAS are observed by the adversary. According to the assumptions, the only information available to the attacker is a *set of message timestamps* (or fingerprints) F.

$$F = R \cup E \cup D \tag{1}$$

where

- R is a set of *regular messages*. These can be from automation rules or reactions to environmental events (e.g. temperature changes) and are of no particular interest to the attacker.
- E is a set of *interesting events* such as direct user interaction (e.g. pressing a light switch) or anything that is of particular interest to the attacker.
- D is a set of *dummy messages* carrying no information.

Messages from other use cases can be put into either group, depending on the scenario. Events from remote user interaction (using an internet gateway) can e.g. be put into R as they leak no information about user presence.

Any of the subsets can be empty. If the HAS consists of a single actuator with a remote control (i.e. no automation rules), then $R = \emptyset$. If it consists of a single temperature sensor periodically sending data to a base station, then $E = \emptyset$.

We assume that $R \cap E \cap D = \emptyset$. If the system supports multiple concurrent channels and the channel does not leak information about a message's contents, the timestamps can be annotated with the channel ID.

3.3 Attacker Model

For a given capture interval $[x, y]$ (x and y are timestamps), the attacker observes the HAS's output $f^{x,y} \subseteq F$. If $t(m)$ denotes the timestamp of message m,

$$f^{x,y} = \{m \in F | t(m) \geq x \wedge t(m) \leq y\} \tag{2}$$

We define all possible subsets $f \subseteq F$ satisfying this condition as the *Subsequence Set* $\mathbb{S}(F)$. Note that when modelling HASs, message timestamps follow a random distribution, so $\mathbb{S}(F)$ is a set of random distributions (one for every possible capture interval) rather than a set of sequences (or sets of concrete timestamps).

4 Privacy Goals

Toledo et. al. have presented a model for information leakage in Private Information Retrieval settings. [17] We formulate our definitions similar to theirs in order to facilitate the development and analysis of countermeasures.

The majority of [17] is not applicable to our setting. In particular, we do not have equivalents for (corrupted) databases/servers. We instead focus on the user (U_t in [17]) and the adversary (A). Furthermore, our Subsequence Set $\mathbb{S}(F)$ relates to the adversarial observation space (Ω): For a given time frame $[x, y]$, the observation space $\Omega_{x,y}$ is the distribution $f^{x,y} \in \mathbb{S}(F)$. An observation ($O$) is a sample from this random distribution. Corresponding to the queries (Q_i, Q_j), the attacker in the HAS scenario provides the user with two *tasks* T_i, T_j (e.g. "Interact with the system during the time $[x, y]$.") of which the user randomly chooses one to execute. The attacker then captures the HAS's traffic (obtaining a sample from the random distribution $f^{x,y}$) and tries to identify the task.

We can thus formulate a notion of privacy in Home Automation Systems.

Definition 1 *A Home Automation System provides (ε-δ)-private communication if there are constants $\varepsilon \geq 0$ and $0 \leq \delta < 1$, such that for any possible adversary-provided tasks T_i, T_j and for all possible adversarial observations O (being a particular random sample of the distribution $f^{x,y} \in \mathbb{S}(F)$) we have that*

$$Pr(O|T_i) \leq e^{\varepsilon} \cdot Pr(O|T_j) + \delta$$

We assume that timestamps are discrete. If they are continuous and $f^{x,y} \in \mathbb{S}(F)$ is a density function, the same definition can be used by substituting the probability for the value of the density function.

As in [17], if $\delta = 0$ we call the stronger property ε-private communication. Note that we require $\delta < 1$. This only affects some cases where a particular observation is certain for one task and impossible for another and prevents the definition from being overly broad.

4.1 Indistinguishability and Unobservability

In practice, if the tasks can be arbitrary, the attacker may choose them to be e.g. "Press the light switch for 10 times in 2 seconds." and "Do not interact with the system for 10 minutes.", producing distinguishable patterns. In order to account for this, we define a slightly weaker property.

Definition 2 *A Home Automation System provides (ε-δ)-indistinguishability for a set of tasks $\mathbb{T}$ if there are constants $\varepsilon \geq 0$ and $0 \leq \delta < 1$, such that*

for all tasks $T_i, T_j \in \mathbb{T}$ and for all possible adversarial observations O (being a particular random sample of the distribution $f^{x,y} \in \mathbb{S}(F)$) we have that

$$Pr(O|T_i) \leq e^{\varepsilon} \cdot Pr(O|T_j) + \delta$$

Probability density functions can also be used here for continuous timestamps. If $\delta = 0$, we call the stronger property ε-indistinguishability.

(ε-δ)-indistinguishability is only defined for a limited set of tasks $\mathbb{T}$. Some tasks are theoretically possible, but unlikely to be encountered in practice. For example, a system might be able to provide unobservability of the user pressing a light switch twice within 10 minutes by making sure that there are always at least two messages in every 10-minute interval. While this does not fulfil the goal of (ε-δ)-private communication, it covers much of the everyday activity and might be more energy efficient than a system offering full (ε-δ)-private communication.

When considering real-world attack scenarios like the detection of user presence [14], the tasks provided by the attacker follow a particular pattern. Instead of choosing two unrelated tasks, the attacker wants to extract a certain piece of binary information from the captured data (such as "Did the user interact with the system?"). In this case, the tasks T_i and T_j from the definition are complementary: $T_j = \bar{T}_i$ (i.e. if T_i is "Interact with the system.", then $T_j = \bar{T}_i$ is "Do not interact with the system."). Due to this being an important special case of (ε-δ)-indistinguishability, we define a separate property:

Definition 3 *A HAS provides (ε-δ)-unobservability of a set of tasks $\mathbb{T}$ if*

$$\forall T \in \mathbb{T} : \bar{T} \in \mathbb{T}$$

and the system provides (ε-δ)-indistinguishability for $\mathbb{T}$.

If $\delta = 0$ we call the stronger property ε-unobservability.

These definitions capture our models as well as real attacks [12,14,13]. Using existing models of user behaviour, one can prove privacy guarantees of a dummy traffic generation scheme.

5 Examples

For trivial approaches it is easy to see whether or not they fulfil the privacy goals.

5.1 No Dummy Traffic

Möllers et. al. have analysed a system which does not produce dummy traffic at all. [14] They have shown that the system does not offer ε-unobservability for the tasks "Interact with the system during a one-hour period." and "Do not interact with the system for one hour." if the attacker knows certain thresholds.

In their experiment, the attacker was able to determine conditions which, if met by the adversarial observation O, would reliably indicate user activity or inactivity. If the predicates $P(O)$ and $A(O)$ denote these conditions, then

$$\forall O : P(O) \Rightarrow$$
$$Pr(O|\text{“Interact with the system”}) > 0 \ \wedge \ Pr(O|\text{“Do not interact”}) = 0$$
$$\forall O : A(O) \Rightarrow$$
$$Pr(O|\text{“Interact with the system”}) = 0 \ \wedge \ Pr(O|\text{“Do not interact”}) > 0$$

As $\nexists \varepsilon : e^{\varepsilon} \cdot 0 > 0$, the system does not offer ε-unobservability, ε-indistinguishability or ε-private communication in general.

In their experiment, the probability of obtaining an adversarial observation meeting the condition if the user performed the given task was less than 1. Thus, the system *may* offer (ε-δ)-unobservability.

5.2 Constant-Rate (Dummy) Traffic

Next, we analyse the concept of *Constant Rate (Dummy) Traffic*. We assume that the system is generating dummy traffic if (and only if) there are no genuine messages to send. Time is divided into slots of fixed length. At the end of every timeslot, either one genuine or one dummy message is transmitted.

Formally, if M is the set containing the ending time of each timeslot and messages in R and E are delayed so that they only occur at the end of a timeslot ($R \subseteq M$, $E \subseteq M$), then dummy traffic is generated by the system so that $D = M \setminus (R \cup E)$.

By construction, the output of the system $F = R \cup E \cup D = M$ is exactly the same, no matter how the timestamps of genuine messages in R and E are distributed. Thus, for any interval $[x, y]$ the adversarial observation will be $O = M \cap [x, y]$, which is stochastically independent from the distributions of R and E. Consequently, for any task T to be executed by the user, it holds that $Pr(O|T) = Pr(O) = 1$. In conclusion, a system using constant-rate traffic provides (ε-δ)-private communication with $\varepsilon = \delta = 0$ (or (0-0)-private communication).

In practice, using Constant Rate Traffic poses a problem. In order to keep the delay for user interaction reasonably low, the overall traffic rate must be very high (i.e. ≥ 1 message per second). However, this can lead to the system violating regulatory thresholds or draining the battery of connected devices. In wired systems, this is generally not an issue. In a wireless setting, other approaches which minimise the generated amount of traffic have to be evaluated. The development of such a system is left for future work.

6 The Legal Framework

As with most topics in the field of security and privacy, technical protective measures and the legal framework for prosecution of offenders are intertwined.

On the one hand, technical countermeasures make attacks more difficult. On the other hand, an effective legal framework can even act as a deterrent to potential offenders. In this context, we highlight problematic areas in legal frameworks using the German Criminal Code as an example. We then present an international effort aiming to improve the legal situation in over 50 countries.

6.1 Legal Challenges

The interception of traffic from HASs (or any other private network, for that matter) can be considered "data theft". However, the definition of theft in the German Criminal Code (Section 242) only applies to *chattels* and not to incorporeal data. [18] Even if the adversary actively intrudes and asserts control of the HAS, the attack does not qualify as trespassing according to Section 123 of the German Criminal Code. The law requires physical entry into a spatially delimited area. [19, § 123, Recital 15]

These two cases show a fundamental problem with many legal frameworks: Criminal law has been developed in times of limited technical prevalence.

6.2 Legal Reforms

In order to update the law and to keep up with new technical developments, legislators have passed reforms. For the German Criminal Code, these are e.g. Sections 202a (Data espionage) and 202b (Data Interception). In our scenario, Section 202b is to be applied. It punishes the illegal interception of data from a non-public data processing facility or from the electromagnetic broadcast of a data processing facility. While it can be argued that air is a broadcast medium and that wireless transmissions are public by nature, the German Criminal Code bases the definition on the intention of the sender.

As criminal law only applies to the respective country, a detailed examination of the Section's contents is outside the scope of this paper. However, Sections 202a and 202b are mere examples of a global development: By changing Section 202a and introducing Section 202b in 2007, the German legislator has implemented the Convention of Cybercrime of the Council of Europe. This guideline, also known as the *Budapest Convention*, has been opened for signature in 2001 and has since been ratified by over 50—European and non-European—countries.

The Budapest Convention is of dogmatic importance for cybercrime and has ramifications on a global scale, especially in the following two areas.

Prosecution Attacks involving computers ("Cyber Attacks") often reach across borders. National solo efforts to combat these are rarely promising. Instead, a coordinated concept supported by as many countries as possible is necessary.

Technical and Legal Terms One part of this coordinated concept is a collection of common terms. In this paper, we assume that the attacker can only access traffic (meta-)data. The distinction between content and traffic (or meta-) data

is important for the legal framework and can be explicitly found in the Budapest Convention. In Article 2 d), the convention states that

> "traffic data" means any computer data relating to a communication by means of a computer system, generated by a computer system that formed a part in the chain of communication, indicating the communication's origin, destination, route, time, date, size, duration, or type of underlying service.

Establishing common terms and ensuring consistent usage is not a trivial task. For example, the German Telecommunications Act considers location data for mobile devices to be traffic data, but this classification cannot be found in the Budapest Convention.

It remains an open question whether the term *data* as used in Article 3 of the Budapest Convention refers to both traffic and content data or only to content data. Answering this question as well as problems resulting from this are beyond the scope of this paper.[4]

6.3 Summary of the Legal Situation

The Budapest Convention contains regulations about criminal law. Certain actions such as the illegal interception of data are to be penalised. This reveals a central idea: Criminal law can be considered part of a comprehensive data protection concept which aims to optimise both technical and legal aspects.

This understanding does not only hold for the German Criminal Code we used as an example here. Instead, it holds for all countries which have signed or ratified the treaty—implementing it in national law. The list of signatures[5] does not only include major European countries such as France, the UK and Italy. It also features many others such as Australia, Canada, Japan or the USA. The questions sketched here regarding the relationship of technology (in this context, especially Home Automation Technology) and criminal law are therefore of international importance.

7 Conclusion

We have established a model for traffic analysis attacks in Home Automation Systems, keeping assumptions general and adapting existing definitions. The model is suitable for developing dummy traffic generation schemes not only for Home Automation Systems, but for networks with similar properties as well. The definitions ensure that privacy guarantees can be mathematically proven.

We have also shown how technology and attacks using it have forced legal reforms and new laws that surpass national borders. The Budapest Convention

[4] This especially holds for questions regarding data retention.

[5] `https://www.coe.int/en/web/conventions/full-list/-/conventions/treaty/185/signatures`, last accessed 10 July 2017.

serves as an important step towards an internationally agreed understanding of terms and necessary actions. While discrepancies between technology and the laws governing it are unavoidable, we have shown that the two can work together towards the goal of protecting people's privacy.

References

1. Chan, H., Perrig, A.: Security and privacy in sensor networks. Computer **36**(10) (2003) 103–105
2. Conti, M., Willemsen, J., Crispo, B.: Providing Source Location Privacy in Wireless Sensor Networks: A Survey. IEEE Communications Surveys & Tutorials **15**(3) (2013) 1238–1280
3. Matos, A., Aguiar, R.L., Girao, J., Armknecht, F.: Toward dependable networking: secure location and privacy at the link layer. IEEE Wireless Communications **15**(5) (2008) 30–36
4. Yang, Y., Shao, M., Zhu, S., Urgaonkar, B., Cao, G.: Towards Event Source Unobservability with Minimum Network Traffic in Sensor Networks. In: Proc. WiSec '08, ACM 77–88
5. Shao, M., Yang, Y., Zhu, S., Cao, G.: Towards Statistically Strong Source Anonymity for Sensor Networks. ACM TOSN **9**(3) (2008) 34:1–34:23
6. Pfitzmann, A., Pfitzmann, B., Waidner, M.: ISDN-Mixes: Untraceable Communication with Very Small Bandwidth Overhead. In: Proc. GI/ITG-Fachtagung '91, Springer 451–463
7. Levine, B.N., Reiter, M.K., Wang, C., Wright, M.: Timing Attacks in Low-Latency Mix Systems. In Juels, A., ed.: FC '04 Revised Papers, Springer 251–265
8. Shmatikov, V., Wang, M.H.: Timing Analysis in Low-Latency Mix Networks: Attacks and Defenses. In: Proc. ESORICS '06, Springer 18–33
9. Dwork, C.: Differential Privacy. In: Proc. ICALP '06, Part II, Springer 1–12
10. Dwork, C., Naor, M., Pitassi, T., Rothblum, G.N.: Differential Privacy Under Continual Observation. In: Proc. ACM STOC '10, ACM 715–724
11. Bash, B.A., Goeckel, D., Guha, S., Towsley, D.: Hiding Information in Noise: Fundamental Limits of Covert Wireless Communication. IEEE Communications Magazine **53**(12) (2015) 26–31
12. Möllers, F., Seitz, S., Hellmann, A., Sorge, C.: Extrapolation and Prediction of User Behaviour from Wireless Home Automation Communication. In: Proc. WiSec '14, ACM 195–200
13. Mundt, T., Dähn, A., Glock, H.W.: Forensic analysis of home automation systems. In: HotPETs 2014
14. Möllers, F., Sorge, C.: Deducing User Presence from Inter-Message Intervals in Home Automation Systems. In: Proc. IFIP SEC '16, Springer 369–383
15. Bergstrom, P., Driscoll, K., Kimball, J.: Making home automation communications secure. Computer **34**(10) (2001) 50–56
16. Bratus, S., Cornelius, C., Kotz, D., Peebles, D.: Active behavioral fingerprinting of wireless devices. In: Proc. WiSec '08, ACM 56–61
17. Toledo, R.R., Danezis, G., Goldberg, I.: Lower-Cost ε-Private Information Retrieval. Proceedings on Privacy Enhancing Technologies **2016**(4) 184–201
18. Vogelgesang, S.: Datenspeicherung in modernen Fahrzeugen – wem „gehören" die im Fahrzeug gespeicherten Daten? juris – Die Monatszeitschrift **3**(1) (2016) 2–8
19. Fischer, T.: Strafgesetzbuch: StGB. 64 edn. C.H.BECK (2017)