PETs in Your Home – How Smart is That ?

HUPS Position Paper

Stefan Korff
University of Münster
Department of Information Systems
Leonardo-Campus 3, 48149 Münster, Germany
stefan.korff@uni-muenster.de

ABSTRACT

Information technology has reached a level of sophistication so that its users leave detailed traces of the parts of their lives when they knowingly interact with information systems (e.g., Internet use) or participate in the social sphere (e.g., video surveillance). The recent advent of the smart home technology, that is residential properties being equipped with sensors and interconnected smart devices, expands the realm of unavoidable data collection further. This exacerbates people's fear of privacy breaches. The purpose of this paper is to evaluate how reasonably the known privacy-enhancing technologies can be applied in the case of a fictional smart home scenario. It briefly discusses potential areas of privacy problems specific to this new technology, then recalls the fundamental ideas behind known privacy-enhancing technologies, and critically evaluates their applicability in an intelligent home with special emphasis on usability. The paper concludes with a discussion of ways forward to effectively adapt privacy protection concepts in this particular environment.

General Terms

Human Factors

Keywords

Smart home, Privacy

1. INTRODUCTION

In a recent issue of Wired magazine, Bill Wasik predicts that we are at the advent of an era where the "most mundane items in our lives can talk wirelessly among themselves, performing tasks on command, giving us data we've never had before" [20]. The so-called Internet of Things, where lights and window shades in your home are adjusted by a central control unit according to your preferences; where your kitchen wakes you in the morning with a scent of freshly grained coffee, and if you are sick you'll get a chamomile infusion. The sensor evolution, where every device is equipped with a processor and is able to communicate with its peers over a near-field communication (NFC) or Bluetooth interface. A translator is not needed, they all speak the same language—the language of the future.

For technology enthusiasts this sounds like a pretty fascinating world, but, as Wasik adds, it also seems like a "scary encroachment of technology". The idea of smart homes equipped with sensor means that the collection of personal data is no longer solely done if you participate in a survey or while you are surfing the Internet. It happens when you are cooking, doing your workout, switch the light off, or even when you measure your current heartbeat. This data is sufficient to create a digital copy of your life. This can be a very unpleasant situation, even for technology enthusiasts.

The Wired article leaves several questions unanswered. What happens to all the data those devices record and generate? Is a smart refrigerator able to decide which information is sensitive or confidential and should not be communicated over a public network? If all devices can communicate with each other, who decides what they are allowed to talk about? From a privacy point of view, these are very important questions. Certainly, machines are not nosy to learn confidential or personal data, but other human beings are. In the case of the refrigerator, your food supplier would probably like to know which recipe we like the most and how much we are willing to pay for certain ingredients. Criminals who seek targets to rob a home could evaluate sensor data to check if someone is at home. At the time of writing, news reports claim that the US government systematically accesses the databases of all major firms whom millions of citizens entrust their data about their personal lives. Imagine such programs being extended to data collected in smart homes.

To curb the uncontrolled collection and misuse of personal data, lawmakers enacted data protection laws. But legal protection alone is not enough. Data protection laws contain loopholes, they consistently lag behind technological development, and effective enforcement of provisions regulating information flow control is somewhere between extremely difficult and impossible.

An alternative approach to protect personal data is the use of privacy-enhancing technologies (PETs). Conceived by engineers and propagated by privacy activists for long, uptake in practice has been disappointing for various reasons including a lack of incentives by the parties in power to define standards coupled with unsatisfactory user experience of typical PETs. Moreover, most of these technologies were designed with conventional ICT in mind: a proficient user interacting with a single, own, and trustworthy device

Symposium on Usable Privacy and Security (SOUPS) 2013, July 24–26, 2013, Newcastle, UK.

connected to the wilderness of the Internet. The purpose of this position paper is to revisit the ideas behind known privacy-enhancing technologies and to evaluate how applicable they would be in a fictional smart home scenario. A special emphasis is put on the usability aspects of adopting PETs for the smart home.

The remainder of this position paper is organized as follows: Section 2 gives an overview about the basic terminology and introduces a fictional smart home scenario which serves as the basis for our later discussions. Afterwards we describe the technological characteristics of smart homes and emphasize the inherent privacy and security problems. Possible privacy enhancing technologies (PETs) which are designed to prevent privacy breaches are briefly described in section 3. Beside a short description of their basic functionality this section strives to give an answer on how usable common PETs are in the particular case of smart homes.

2. PRELIMINARIES

2.1 Terminology

The term "smart home" is not precisely defined. Some authors describe a smart home as a holistic home healthcare system, which facilitates elderly or disabled persons to live more independently in their own residence [8]. Others subsume as broad range of technologies which are dedicated to interconnect and control devices in our houses with the purpose of providing greater convenience, safety, security, and energy savings to the residents [21]. Our working definition of a smart home is the concept of digitally interconnecting and controlling devices, which are primarily designed to be used in a home (such as refrigerators, televisions, toasters, window shades, door locks, etc.).

Further, we assume that the data communicated in the smart home network can be rather general "temperature +1 degree" but also very intimate "start to order baby food in 9 months".

The system itself is believed to provide a set of rudimentary security mechanisms. Those are useful to protect the home network against a less complex class of attacks. But as in the case of most networks, there are various loopholes which can be exploited by a more sophisticated attacker. Therefore, we presume that for instance an eavesdropping operation is one of many realistic adversary scenarios we need to consider. In the following section we outline some reasons why these scenarios are not unrealistic and why the protection of privacy should be taken serious.

2.2 Sensors and Devices

The composition of interconnected nodes in a smart home network encompasses a wide range of technical devices and gadgets. Those can be kitchen devices, smart meters or domestic health care systems. The kind and source of data which are supposed to be recorded can vary between the systems. Probably the most intimate data are collected by healthcare devices. Even those systems have security and privacy weaknesses, as demonstrated in [1] and [18]. The attack scenarios described in these projects range from unauthenticated access, via message modification, to denial of service attacks. A reason for this vulnerability might be that devices designed for smart homes are not always developed with an adversary scenario in mind. Moreover, the vendors may have an interest in collecting and analyzing the user data to improve their products. In the case of the medical devices, it might be in the user's interest to transmit the test results to their responsible doctor. However, this increases the risk that the data is forwarded to third parties who are, for instance, interested in the household's consumer behaviors.

2.3 Network and Inter-Connectivity

A smart home consists of different interconnected devices. Some are primarily communicating with a central control unit while others establish multiple communication links to internal and external counterparts. In order to transmit information between the devices, developers proposed various network communication protocols. One widely adopted open source protocol is KNX [1]. KNX is standardized by the ISO and administered by the KNX Association which currently has more than 300 members, majorly technology companies. The protocol is platform independent and can be used to transmit messages over a wired connection like Ethernet or Powerline, or wireless by using infrared radiation. Further, there is a modification called KNXnet/IP which is using IP networks as a KNX medium.

Various protocols like KNX have their origin in a time where a closed centralized architecture was the predominantly used system design. However, the closed system assumption was probably one reason that the implementation of effective security mechanisms has been neglected.

For instance KNX only provides as basic access control management but no mechanisms to prevent network attacks [13]. KNX/IP admittedly offers a richer set of security mechanisms and uses encryption technology like AES and HMACs [16]. However as noted by Granzer et al. some of the used mechanisms are still violating Kerckhoffs principles [15] by following a "security by obscurity" approach [13]. The same authors exposing multiple other vulnerabilities in various building automation system protocols. We admit that the presented scenarios have been investigated in the year 2010 and major vulnerabilities haven been probably closed. Despite, it shows that these technologies are not immune against malicious behavior.

2.4 Business Strategies

Obviously developing smart home technology is not a philanthropic idea of the industry, it's a business. Even if a smart home might improve the living standards and can be designed to help elderly people, it is still a source of revenue. The strategic alignment of a company can indirectly influence the design of PETs and their usability. For instance, from a PET and usability perspective it might be desirable to have standardized, open and decentralized system architectures. However, due to a new business strategy a hardware vendor is building a new heat sensor which is designed to synchronize data with a centralized platform. The same company is building its business on a customer-lock in strategy: A strategy where the vendor is establishing high switching costs to prevent a user from changing to competitive technologies [3]. This strategy usually implies a closed and proprietary system design. Both scenarios are contradicting the initially favored open system design. This shows that not solely technological but also business interests can influence the effectiveness and usability of PETs.

[1] KNX: http://www.knx.org/knx-standard/introduction/

3. USABLE PETS FOR SMART HOMES

Safeguarding privacy in a smart home environment is a challenging task because there is

1. a multitude of devices which collect or generate data,
2. the devices are diverse in technology and build on different standards; so getting the functional properties work may require some effort, which makes it less likely that people care about non-functional properties, such as security and privacy features;
3. the amount of data is huge,
4. most devices are connected to an (ad-hoc) network,
5. for cost and convenience reasons, few devices will have a rich user interface with screen and controls needed to configure security and privacy settings and to identify legitimate users.

In this paper, we are not giving solutions for these problems, but rather try to scrutinize which of the fundamental privacy enhancing technologies are applicable in smart homes. We shortly introduce each technique and reveal possible problems concerning their adaption.

3.1 Data Minimization

The strongest technique to prevent privacy breaches is to avoid the recording of personal data. This might be a pretty simple requirement but it is only rarely adapted in information systems. For these classes of systems which functionality is related on the input of personal data, the next requirement would be not to store the data. If they necessarily have to be stored the system must provide the user with the opportunity to deactivate this functionality.

To shorten the discussion at this point, a system which is not recording personal data contradicts to the very basic idea of smart homes. They are built to collect data, transfer them and use them for decision making. There might be some devices which necessarily don't have to store them but still they are applied to record them. Despite, we can relax the stringent requirement on complete data avoidance and restrict the recording to those data which are mandatory in order to enable the actual functionality of the device. For instance, a smart refrigerator can be used for bookkeeping the storing and withdrawing of groceries but should not track if someone stores a medicine which needs to be kept on a low temperature.

3.2 Dummy Traffic

The concept of dummy traffic was proposed as a countermeasure against intersection attacks [5]. These attacks aim to deduce information from communication patterns and to uncover the users anonymity provided by a used anonymity service. The analysis requires a short to long term observation of the communication lines in order to learn something about their activity. If a user switches inactive the attacker will get additional information which helps him to link the message pattern to a certain user. Dummy traffic is countering this attempt by simulating activity even if the real sender is actually not sending any message. By concealing the inactivity of a user the attacker is detained from learning additional information about the identity of the users. In the case of a smart home the generation of dummy traffic is possible but hardly practical. For instance, a user intends to prevent the refrigerator from learning something about his food preferences and to use them as an identification pattern. Obviously for a user it requires more effort to conceal groceries than to decrypt a digital message with a software tool. He can either unpack all things he bought and use a unique packaging or only buy things he does not like and his preferences stay confidential, as illustrated by in figure 1. Both options are not typical examples of perfect usability.

„In the smart home, you have to eat Whiskas
so that noboby knows you're a dog."
©The New Yorker Collection 1993 Peter Steiner
From cartoonbank.com. All rights reserved.

Figure 1: Dummy traffic in the case of smart homes. A modified version of Peter Steiner's famous cartoon published by The New Yorker on July 5, 1993.

3.3 Differential Privacy

Another concept which attempts to confuse the adversary by manipulating an observed set of data is differential privacy. The purpose of this approach is to provide a rigorous protection of all personal identifiable information of an individual in a statistical data set [10]. The adversary scenario which is addressed by this approach is asserting that by knowing a result R aggregated from a set of statistical data D_I an attacker retrieves information about an individual. For instance R can be the aggregated query result of a data set $Q(D_I)$ which was created by empirical data assessment. By knowing R an attacker automatically can learn new information about an individual person. This is invariant no matter if the persons contributed something to the set of answers D_I or not. Further, knowing R the attacker can estimate if you participated in the survey or not. Minimizing the probability of finding out if you participated in the survey or not is the actual focus of differential privacy. To minimize the evidence if an individual contributed his data following equation must hold:

$$\frac{P(R \mid Q(D_I))}{P(R \mid Q(D_{I\pm i}))} \leq e^{\epsilon} \text{ for all } I, i, R \tag{1}$$

According to equation 1 the chance that the released results would be R is nearly the same whether an individual participated to the survey D_{I+1} or not D_{I-1}. To satisfy this requirement the global sensitivity ΔF of the sets $D_{I\pm i}$ is calculated. The global sensitivity is measuring the difference caused by adding or removing an element from the data set. To satisfy equation 1 this gap needs to be spanned. This can be realized by adding random values generated by a Laplacian distribution to the data distribution. The equalization of the sets D_{I+1} and D_{I-1} by adding noise is resulting in the privatized version of R. Since ΔF is now smoothed out, it gets very hard (e^{ϵ} hard) for adversary to estimate if an individual participated in the data set by just knowing R.

In the case of a closed smart home environment, differential privacy has at least three problems. First, it is a beautiful theoretical concept, but if the database D consists of many variables with relatively few independent records, any reasonable choice of ϵ implies that the query result is blended with so much randomness that its utility for control purposes barely deserves the attribute “smart”.
Second, concurrent queries with differential privacy guarantees require a central entity to keep track of the query history and random values. Otherwise attackers can improve their inference by launching concurrent requests. Smart homes implementing an ad-hoc network architecture lack this central entity.
Third, the randomization impairs the user experience. How can you explain your grandma a position on her electricity bill saying “Random charge to protect your privacy”?

3.4 Informed Consent

A common way to set-up agreements in information systems is the use of consent dialogues. In an informed consent the user is instructed about the disclosure of personal data as well as other regulations like an end-user license agreement (EULA). There has been extensive research on the optimal layout of consent dialogues in the past [12]. The research strives to design a comprehensive presentation of the terms and condition without influencing the users decision on accepting or declining the agreement.
Despite the incentive to provide the user with a transparent disclosure, research have also shown that users are bothered by interception dialogues and “trained to accept” them [11]/[7]. Accepting an agreement is rather seen as a mandatory and not an optional choice. Moreover studies have shown that users tend to ignore the different content and purposes of a displayed dialogues. It doesn’t matter if the presented consent is a privacy agreement or a EULA [7].

The habituation of ignoring and the unreflected acceptance of agreements and policies will also be present in the case of smart home technology. However, devices which have a limited functionality are presumably configured a single time. Those devices can be configured by a central policy as described in section 3.5. Further, some devices are only incidentally noticed and cannot trigger an active dialog with the user. If a scenario emerges which was previously not covered by the preconfigured privacy policy the user must be aware of that.
In the case of rather actively used devices, other problematic scenarios might emerge. In case of the smart refrigerator example following situation might be possible: After a software update, which added some additional sensor functionality, the device rolls out a new privacy policy. If the user wants to withdraw food, he first needs do accept the new policy. In the case he is very hungry, he will be bothered to read all the new conditions and simply clicks on accept. To prevent that, the device can give him the possibility do postpone this decision. The new functionality will be deactivated till he is explicitly accepting the resulting changes of the privacy policy. Another case would be if the user asks for a replenishment of groceries. The device offers a selection of online shops where the user can place his order. Whereas every shop has a different privacy regulation. By using an integrated touch screen the user has the opportunity to browse through the policies and accept or refuse the terms and conditions. In this case we again might run into the mentioned problem of a prompt acceptance. He doesn’t want to read policies he wants to get his refrigerator refilled. A remedy in this situation would be if the system compares a predefined policy with the policies of the online merchants and sorts out the mismatches. The user is no longer bothered by reading notifications and can concentrate on his main task.

3.5 Policy Languages

If the avoidance of personal data recording is not an option a company needs to consider dedicated legal regulations. Those are requiring a company to transparently display the type and intended usage of the captured data. The information is incorporated into a legal document, the privacy policy.
As described in the previous section a privacy policy must be accepted by the user. However, we learned that even if this decision has consequences for the confidentiality of personal data many users are accepting them immediately without any doubts. To counteract this trend various developers and researcher came up with the idea of providing designers of information systems with a framework for the embedding of privacy policies. The purpose of these frameworks is to enable a more transparent, readable and usable presentation of privacy policies.
Essential elements of these frameworks are policy languages. Those can be understood as dedicated markup languages for privacy policies. The Privacy Preferences Project (P3P) was developed by the W3C as a machine-readable and actionable policy language. The policy which is written in this language can be interpreted and directly displayed by the user’s browser agent. Moreover the user has the possibility to predefine its own privacy policies. If he encounters a website which is following a privacy-by-policy approach his personal privacy settings are automatically matched against the privacy policy of the website and differences are reported [9]. It has been demonstrated that P3P can also be used to generate a human readable version of the privacy policy by using a fixed taxonomy. The existence of a disclosure will only be helpful to a certain extend. More important for the understanding of the user is an appropriate usable and intelligible policy [9]. Therefore, the presentation of a policy should be always adapted to the systems environment.

The usability of a privacy policy as mentioned is depending on the systems architecture. So far languages like P3P are written for web browser extensions which make them more or less platform independent. But at least they require some kind of output and input device. A problem which needs to be faced here is that devices in smart homes might have only a very limited screen size or even no screen at all. It will be a challenging task to insert privacy preferences into a clinical thermometer or a smart toaster.

In terms of usability it makes sense, to provide those devices with a centralized platform where a user can manage his privacy preferences for each device connected to the network. To avoid heterogeneity of policy layouts and definitions it would be helpful to implement a common naming and description of data by using a fixed taxonomy. As mentioned, there are various network communication protocols and systems available for connecting a smart home. Therefore, it might be appropriate to implement a generic policy framework. Otherwise a toaster cannot connect to the network due to a non-interpretable privacy policy.

3.6 Identity Management

A common concept of modeling different areas for privacy is the use of sphere models. Spiekermann and Cranor [19] use such a model to explain different domains of responsibility for privacy protection whereas others are using spheres to classify different areas of sensitivity [14]. So far we did not consider that a home and the accommodated household can consist of more than one member. Each member no matter if they are part of a family, living community or a partnership has its own sphere of privacy. In order to protect the individual needs for privacy it is important that systems which record private data are aware of these different identities. Especially if those systems can be used by any member of the household.

To limit the scope of the discussion we focus on the granule of identity management, the identification function. In the case of a computer system a common practice is to create a person related user. The person can then be identified by entering the matching password token into the system. The purpose of identity and access management is to restrict the access to personal data stored by the system. After an individual identity has been verified he can access the personal data. The field of identity and access management comprises a multitude of mechanisms which are also possible candidates for an enrollment in smart home environment.

In the case of smart home devices the entering of a password token might be problematic. As mentioned some devices don't have any screen and can only provide a very limited amount of input values. However, identification can play a crucial role and needs to be managed. In the case of domestic health care devices which are syncing data with a medical database, a missing identification mechanism would wreck the whole functionality. One authentication scheme proposed for smart healthcare systems is using physiological information detected by a biomedical sensor to identify a person [4]. Such a mechanism would be also a feasible solution for a smart home. However from a security perspective the exchange of secret credentials with every device is problematic. Imagine someone buys a used smart toaster at a garage sale and afterwards naively integrates the new device into his home network without any further checks. What he doesn't know, is that the device is infected with malware. The now actual "physical malware" is designed to collect the user's credentials and to transfer them to a potential attacker. To avoid such a procedure, the system can use an approved single sign on solution similar to Kerberos. By using such a technique the user needs to initially exchange the credential with an authentication device (trusted key server). Afterwards, a ticket is issued which can be used for requesting access to all other devices without any needed additional exchange of credentials.

Another technique which might be a feasible solution in smart homes is distance bounded authentication. An approach which is used by the wireless communication standard Bluetooth [6]. In this case a user needs to be initially paired with a device. The relation will be memorized and whenever the user is in a certain distance to the device he will be automatically granted access rights. This model actually corresponds to the non-digital regulation of access control. In general everyone who can physically access my toaster and asked for my permission is allowed to use it. Such an ad-hoc regulation of access rights would also prevent the creation of complex access control matrices in case of devices which are shared by multiple users. For the pairing itself we can possibly use a physiological signal based identity authentication.

Nevertheless, in the case of a required remote access it might be necessary to implement a central transparent access control management. Of course such a system should comprise basic security mechanisms like the automatic decay of access rights. For instance, after their graduation students are no longer allowed to live in the dormitory.

4. DISCUSSION AND CONCLUSION

This paper provides a brief discussion on the usability of known privacy enhancing technologies in a smart home. We focused on the fundamental techniques starting with the principle of data minimization and closing with identity management. We considered rather technical approaches like dummy traffic and the more recently discussed concept of differential privacy.

The smart home is a technological movement which final area of deployment is still hard to grasp. Also the willingness to adopt this technology cannot be foreseen yet. The paper demonstrates that it will be rather challenging to use known PETs in a smart home context. Some techniques like "don't collect data at all" are not working by definition. Taking into account the usability perspective also dummy traffic will be hard to implement. Due to the diversity of technology standardized identification mechanisms which enable a cohesive authentication chain are hard to realize. Despite that the findings are rather disillusioning, we like to discuss some ideas for designing a PET for the smart home.

In the beginning is useful to recapitulate the discussed preconditions: First, it would be naive to fully trust into the security of the devices and networks. We mentioned that many communication standards are not designed by determining security as the primary objective. As in all other networks we need to take potential privacy breaches into account. Second, we cannot ignore different business interests of vendors and industries. Therefore, it is important to design privacy architecture as independent as possible. Finally, we need to incorporate the different users and user groups and their respective combination into our privacy model. The distinctive forms of living and sharing of devices like a classical family model or shared dormitory room have to be taken into account.

To give an idea how the mentioned preconditions can be adopted we describe a possible scenario for privacy policies and identity management.

As it was mentioned in section 3.4 users are prone to ignore warnings and simply agree on every dialog they get confronted with. To encounter this, the vendors can deliver the devices with a very restricted default configuration even if this entails a reduced functionality. This presupposes a dedicated mapping between each functionality and the used personal data. By defining privacy on a functional level (which function uses which data) the collection of non-functional driven data can be reduced. Further, we recommend to install a centralized and transparent privacy control. To achieve a reasonable usability and interoperability, standardized frameworks and policy languages like P3P should be adopted. For a rigorous privacy protection it has to be mandatory for any device to adopt the networks central policy regulation before it can be used.

Attributes as transparency, interoperability and the usability are also basic building blocks for the design of the identity management. The identification mechanisms must be able to flexibly adapt to the above mentioned combinations of users and user groups without generating an unnecessary administration overhead. A first approach would be the use of individual privacy zones as suggested by Arabo et al. [2]. The zones separate the user related data on a shared device. However, a detailed identification process is not described in this framework.

One approach we suggested is the use of a single sign-on solution. A home network authentication protocol inspired by Kerberos can be used to supersede the exchange of credentials with every device in the network [17]. Further, the administration effort can be limited to a single authentication device. Another possible idea is to realize access regulation based on a distance boundary model. This technique corresponds to traditional access control to shared devices within a household.

No matter if and how a current PET is finally applied, it should be considered that they always reduce the usability on the first instance. Therefore, it is challenging to manage the actual tradeoff between the perceived usability and the perceived benefits of privacy enhancing technology. Every notification or dialog will distract the user from carrying out his actual planned task. Therefore, it is important to use the right technique in the right situation to seamlessly integrate the privacy protection mechanisms into the workflow of the user. Finally it will be the task to prevent that the "scary encroachment of technology" becomes a privacy nightmare.

5. REFERENCES

[1] M. Al Ameen, J. Liu, and K. Kwak. Security and privacy issues in wireless sensor networks for healthcare applications. *Journal of medical systems*, 36(1):93–101, 2012.

[2] A. Arabo, I. Brown, and F. El-Moussa. Privacy in the age of mobility and smart devices in smart homes. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)*, pages 819–826. IEEE, 2012.

[3] W. B. Arthur. Competing technologies, increasing returns, and lock-in by historical events. *The economic journal*, 99(394):116–131, 1989.

[4] S.-D. Bao, Y.-T. Zhang, and L.-F. Shen. Physiological signal based entity authentication for body area sensor networks and mobile healthcare systems. In *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, pages 2455–2458. IEEE, 2005.

[5] O. Berthold and H. Langos. Dummy traffic against long term intersection attacks. In *Privacy Enhancing Technologies*, pages 110–128. Springer, 2003.

[6] S. I. G. Bluetooth. Specification of the bluetooth system, version 1.1. *http://www. bluetooth. com*, 2001.

[7] R. Böhme and S. Köpsell. Trained to accept?: a field experiment on consent dialogs. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 2403–2406. ACM, 2010.

[8] M. Chan, D. Estève, C. Escriba, and E. Campo. A review of smart homes—present state and future challenges. *Computer methods and programs in biomedicine*, 91(1):55–81, 2008.

[9] L. Cranor. *Web privacy with P3P*. O'Reilly Media, Inc., 2002.

[10] C. Dwork. Differential privacy. In *Automata, languages and programming*, pages 1–12. Springer, 2006.

[11] S. Egelman, L. F. Cranor, and J. Hong. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1065–1074. ACM, 2008.

[12] B. Friedman, P. Lin, and J. K. Miller. Informed consent by design. *Security and Usability*, pages 495–521, 2005.

[13] W. Granzer, F. Praus, and W. Kastner. Security in building automation systems. *Industrial Electronics, IEEE Transactions on*, 57(11):3622–3630, 2010.

[14] B. E. Hermalin and M. L. Katz. Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quantitative Marketing and Economics*, 4(3):209–239, 2006.

[15] A. Kerckhoffs. *La cryptographie militaire*. University Microfilms, 1978.

[16] D. Lechner, W. Granzer, and W. Kastner. Security for knxnet/ip. In *Konnex Scientific Conference*, 2008.

[17] S. P. Miller, B. C. Neuman, J. I. Schiller, and J. H. Saltzer. Kerberos authentication and authorization system. In *In Project Athena Technical Plan*. Citeseer, 1987.

[18] H. Ng, M. Sim, and C. Tan. Security issues of wireless sensor networks in healthcare applications. *BT Technology Journal*, 24(2):138–144, 2006.

[19] S. Spiekermann and L. F. Cranor. Engineering privacy. *Software Engineering, IEEE Transactions on*, 35(1):67–82, 2009.

[20] B. Wasik. Welcome to the programmable world. *WIRED magazine*, 06-2013:140–147, 2013.

[21] N. Wingfield. Controlling the 'smart home' with tablets and smartphones. *http://tinyurl.com/mb59luu*, 2013.