Consumer Smart Homes: Where We Are and Where We Need to Go

Shrirang Mare, Logan Girvin, Franziska Roesner, Tadayoshi Kohno
 Allen School of Computer Science & Engineering
 University of Washington
 Seattle, USA
 {shri,lsgirvin,franzi,yoshi}@cs.washington.edu

ABSTRACT

Currently a variety of smart home systems are available from different vendors, who made different design choices in how these systems operate, which in turn affect end users. A key question that we consider in this work is how these commercially available smart home systems differ in practice, what are the implications of those differences, and whether other design alternatives might be better. To answer these questions, we systematically evaluate seven popular smart homes and identify their underlying design choices around access control, privacy, and automation, and highlight the implications of those design choices for end users. We surface challenges and tensions for design choices around topics like security, privacy, usability, automation, and reliability, and we make design recommendations where possible. Our findings lay the groundwork for future research in this area.

KEYWORDS

Automation, Internet of Things, Privacy, Security, Smart Home

ACM Reference Format:

Shrirang Mare, Logan Girvin, Franziska Roesner, Tadayoshi Kohno. 2019. Consumer Smart Homes: Where We Are and Where We Need to Go. In *The 20th International Workshop on Mobile Computing Systems and Applications (HotMobile '19)*, February 27–28, 2019, Santa Cruz, CA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3301293.3302371>

1 INTRODUCTION

In recent years, the number and variety of consumer smart home devices has increased rapidly. The mobile and transient nature of smart home computing (e.g., occupants come and go in a home, devices can be controlled via mobile phones when a user is at home or away, guest users come and go but may wish to control the devices when they are home) raises several design challenges around topics such as access control, automation, and privacy [1, 5, 6]. It is unclear what design practices current smart homes follow, how they address the tensions and challenges raised by past research, and what new (if any) challenges and implications (for end users)

these modern smart homes introduce. In this paper we attempt to fill this gap by evaluating seven modern consumer smart homes.

Our intent is not to cast stones at any particular smart home vendor. We seek to illuminate the current design choices in consumer smart homes and surface the many tensions. We hope our findings will be useful to product developers to refine their products, and the tensions that we identify will help guide future research in this space.

Much of the prior work on modern smart home devices has been focused on a particular device, a particular smart home, or a particular component of a smart home (e.g., [7, 15, 22, 23]). In contrast to studying a specific smart home product, we take a broader view and collectively evaluate seven different smart homes to understand the different design alternatives and their implications on end users. We based our evaluation on three prominent themes in smart home computing: access control, privacy, and automation. Along these themes, we seek to answer questions such as: How do different smart homes address access control, accounting for changes in users' mobility and context? How do different smart homes handle privacy within a home? How do different smart homes handle automations?

Contributions. We make three contributions.

- (1) We experimentally evaluate seven different, commercially available consumer smart home platforms, to assess design similarities and differences, and gauge the impact of common and different design decisions on end users.
- (2) Our research surfaces a number of fundamental tensions between different stakeholder values – including security and privacy, usability, and reliability.
- (3) Driven by these findings, we propose technical directions for both industry and the academic community – directions that, if pursued, can help provide a foundation for navigating stakeholder values in the future.

2 EVALUATION METHOD

In our evaluation, we focused on smart home hubs and not on individual smart devices, because hubs serve as platforms on which smart homes are built and they also provide an interface through which users interact and control their smart home. We chose seven smart home hubs to evaluate: Amazon Echo, Apple Home, Google Home, Phillips Hue, SmartThing, Vera, and Wink. To evaluate each smart home hub, we created separate testbeds that included the smart home hub, two sensors (a motion sensor and a door sensor), an actuator (smart power outlet), and a router. We wanted to study each smart home from the perspectives of three types of users: the primary user, occupant user who first sets up the smart home; the

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 the author(s) 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.

HotMobile '19, February 27–28, 2019, Santa Cruz, CA, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6273-3/19/02...\$15.00

<https://doi.org/10.1145/3301293.3302371>

secondary user, other occupant users of the smart home; and the guest user, non-occupant users who may visit the home just for a few hours or for days. So with each home, we set up three mobile phones, one for each type of user.

We evaluated the smart homes in a lab setting for a week as expert end users. Leveraging prior work on how end users use smart homes [22], we systematically designed our evaluation to mimic the different stages of smart home usage, and we created a rubric with instructions for each stage. Specifically, the rubric included instructions for setting up smart home hub, adding and removing devices, adding and removing users, editing the home configuration (e.g., automations to create), and finally, dismantling the home. The rubric also included instructions to evaluate access control mechanisms, test automations, and observe privacy controls available to end users.

Using the rubric, one researcher conducted the experiments with all the seven smart homes, and took detailed notes on the output and state of each smart home during each step in the rubric. Two researchers analyzed these detailed notes to understand the design choices and mechanisms around access control, privacy, and automation in each of the smart homes. We conducted our experiments in May 2017 and again in September 2018 to check whether our initial findings were still valid after one year (we found they were).

Threat model. Any malicious actor in the smart home ecosystem may pose a threat to a smart home user's security, privacy, or safety. For example, a malicious device manufacturer or a malicious smart home app may monitor or control a smart home without the smart home user's consent, or in a multi-occupant home a malicious occupant may spy on other occupants. Some actors outside the smart home ecosystem (e.g., neighbors, guests, remote online adversary) may also compromise a smart home by gaining physical access or remote access to one or more devices in the smart home. In this paper we mainly focus on the security and privacy threats from a malicious smart home occupant to other occupants and guests in the smart home.

3 FINDINGS

We group our findings into the three main themes in smart home computing: access control, privacy, and automation. Table 1 gives a visual summary of our findings.

3.1 Access Control in Smart Homes

Access control is fundamental to any multi-user OS or platform, including smart homes, and so we studied how different smart homes handle access control. Overall, we found that different smart homes have different access control models, which they use to control occupants' access to different types of data and resources in the home; some smart homes attempt to improve usability by offering multiple interfaces to control devices, but this feature can undermine access control in the smart home.

Different access levels. One prominent difference in the smart homes we studied is the types of users (and the granularity of access controls) that they support. For example, on one end of this spectrum is SmartThings, which supports only one type of user as it treats all users in a home equally, giving them the same level

of access and the same view of the smart home, and on the other end is Vera, which supports four types of users, each with varying levels of access: *administrator* (owner privilege), *advance* (administrator user privileges but cannot add/remove users), *basic* (advance user privileges but cannot add/remove automations or scenes), and *notification-only* (can only view notifications). The other smart homes we studied fall somewhere on this spectrum. Wink and Apple Home have two levels of access: *primary user* (who has owner privileges) and *secondary user* (who, by default, can operate devices but cannot add or remove devices or users). Furthermore, in Apple Home and Wink, the primary user can increase the privileges of a secondary user, allowing that user to manage devices and users. Amazon Echo, Google Home, and Philips Hue do not support multiple users. Thus, different smart homes have different approaches for access control.

Different use of access policies. Smart homes use policies to control access to use (or operate) devices, add or remove users and devices, organize devices in scenes or rooms, and add or remove automations. Generally, in all smart homes, the primary user has the highest (owner) privileges, but the privileges of the secondary user differ among smart homes. In Vera and Wink, the secondary user cannot add or remove other users but can add or remove devices and create automations. In SmartThings all users have the same privileges.

Risk of undermining access policies. To improve usability and reliability, some smart devices can be configured such that they can be controlled via multiple interfaces. For example, an outdoor smart camera can be integrated with a voice assistant like Amazon Echo, such that it can be controlled through the smart home hub (the default interface) as well as through the voice assistant (an alternate interface). However, alternate interfaces can undermine a device's access control policy (set by the user) if the alternate interface cannot enforce access control. For instance, in the previous example, if the voice assistant cannot recognize the user issuing voice commands, anyone near the voice assistant can control the outdoor camera, irrespective of the access control policies set for the camera.

Guest access. The transient nature of guests presents a unique access control challenge in smart homes – how to easily grant and control access for guests? As shown in Table 1, different smart homes have different strategies for guest users. In Wink and Apple Home, a guest user can be given access to *certain* devices, whereas in Vera, a guest user can be given *limited* access to *all* the devices. In Apple Home, guest access can be restricted (based on location) to only when the guest is at home. In the other smart homes we tested, once a guest is given access to a smart home device, the guest can control the device remotely and retains access until it is revoked. This approach of managing guest access (where the owner has to grant explicit access and then revoke it) is burdensome for smart home users.

3.2 Privacy

We consider threat to occupants' privacy from external entities and from other occupants in the home.

Privacy from external entities: lack of transparency. A smart home *continuously* collects data about *all* its occupants when they

Property	Ref.	SmartThing	Wink	Apple Home	Vera	Phillips Hue	Amazon Echo	Google Home
Hub category		Multi vendor	Multi vendor	Multi vendor	Multi vendor	Single vendor	Cloud only	Cloud only
Supports multiple users	[18, 24]	●	●	●	●	●	○	○
Access control								
Number of user types		2	2	2	4	1	1	1
Supports guest mode	[10]	○	○	●	○	○	○	○
Supports device level access	[12]	○	●	○	○	○	○	○
Time-based policies	[12]	○	○	○	○	○	○	○
Location-based policies	[3]	○	○	●	○	○	○	○
Privacy								
Detects user presence		●	●	●	●	○	N/A	N/A
User owns the data	[17]	Unclear	Unclear	●	Unclear	Unclear	Unclear	●
User can delete data	[14, 17]	Unclear	Unclear	●	Unclear	Unclear	●	●
Data shared between users	[4, 19]	All	Some	None	Some	None	N/A	N/A
Location where data is stored	[9]	Cloud	Cloud	Cloud	Cloud	Cloud	Cloud	Cloud
Automation								
Communication model		Mixed	Mixed	Mixed	Mixed	Mixed	Mixed	Mixed
Provides-test-environment	[16, 23]	○	○	○	●	○	○	○
Sends device unavailable alert	[5]	○	○	○	○	○	○	○
Network resilience	[6]							
Local processing		●	●	●	●	●	○	○
Local control		○	●	●	●	●	○	○

Table 1: Smart home comparison; *Refs.* column lists example prior work we used to derive the corresponding property.

● indicates the smart home supports the property; ○ indicates it does not.

are home, but also when they are away (through the smartphone app on their mobile phones). Data captured by smart homes can reveal occupants' daily activities and potentially sensitive information about the occupants. We found significant differences in vendors' approaches in handling users' smart home data. As shown in Table 1, Apple Home and Google Home clearly state that data ownership is with users, and they offer users an option to delete their data, but with the other smart homes data ownership policies are unclear.

Data sharing between users. The smart homes we evaluated share three types of data between users: user's *location*, captured through the user's smartphone app (if enabled by the user); *automations* created by the user; and *device activity* log. For these data types, different smart homes offer sharing preferences. In SmartThings everything (*location*, *automations*, and *device activity*) is shared with all users; in Vera *automation* and *device activity* are shared with all users; in Wink only *device activity* is shared; and in Apple Home only *automations* are shared between users (and users can edit each others automations). Furthermore, in Vera, users can see other users' *automation*, but the details vary depending on the user's access level.

Side-channel privacy leaks. Device activity log can reveal a user's habits and behavior that the other users in the home may

not otherwise notice. A smart home's decision on whether to share this information readily between users determines how easy (or difficult) it will be for a malicious user to track another user, and this, in turn, may influence how people use smart homes. Without easy access to this information through a smart home, spying on other users requires intent, skill, and effort; with smart home readily providing this information, spying on other users requires only intent. In addition to device activity logs, unintended private data sharing – privacy leaks – can happen in other ways. For example, in SmartThings, we found that a user can determine another user's presence at any arbitrary location by misusing the user-presence feature. SmartThings shows a user's presence (e.g., at home or away) based on the location of the user's phone and the hub's location. But any user in the home can change the hub's location to any location of interest to check whether other users are present at that location (e.g., is Bob really at work?).

3.3 Automation

For smart homes to reliably execute automations, they have to be aware of the state of the devices in the home and they have to be resilient to network failures [5].

Home state awareness. The communication model between the hub and the devices determine how quickly a hub would learn of any

change in its devices. We considered three communication models: *pull* (hub periodically polls devices for their status), *push* (devices push their status when there is a change), and *mixed* (combination of push and pull). An aggressive polling approach can capture changes in device status quickly, but this approach is expensive in terms of energy, especially if we expect the battery operated smart devices to last for years. A push model is energy efficient but may lead to an inconsistent view of a home if there are network errors and the delivery of certain push messages fails. A mixed model can leverage the energy efficiency of the push model while using a relaxed pull strategy to maintain a consistent view of the home over time. We found that all smart homes have converged to the mixed communication model.

Network resilience. Much of a smart home’s utility is due to it being connected to the internet, but network outages do occur, and therefore, a smart home should be resilient to network failures. We argue that in the event of a network outage smart homes should provide a minimum level of service, which includes supporting devices and services necessary to keep the home in a safe state. For instance, even in the event of a network outage, occupants should be able to turn on smart lights, a smart thermostat should operate normally and not accidentally leave the home in an undesired state – as it happened during the recent outage due to the widespread DDoS attack [20] – and smart door locks should still operate so people can enter or leave the home. We therefore tested smart homes for their support for *local processing*, i.e., whether the smart home can process automations and communicate with devices without internet access, and *local control*, i.e., whether a user can control the smart home without internet access. Overall, we found that smart homes are converging towards a local processing and local control design – a good design choice for a reliable smart home experience, even in the event of a network outage.

Broken automations. Automations can fail because of several reasons such as malfunction of the input (trigger) device, malfunction of the output (actuator) device, or a network failure. We tested how smart homes handle failures in automations due to the unavailability (malfunction) of the trigger device or the actuator device. Specifically, we investigated whether a smart home sends an alert to users when a device becomes unavailable, which is important because the failed device could be a trigger or an actuator for an important automation. In our experiments, none of the smart homes alerted their users if a device became unavailable (Table 1). When a device was unavailable, the device appeared as “unavailable” or “inactive” on the status dashboard of all the smart homes, but users were not notified of the device failure or the affected automation. We also tested how smart homes report execution of an automation when the automation is triggered but the actuator device is unavailable to complete the automation. In Vera, if the actuator was not responsive, the automation status was shown as failed, but in SmartThings and Philips Hue the automation was shown (incorrectly) as successful. Apple Home did not show status or log of previously run automation, so users would not be aware of the failed automation.

4 DISCUSSION

Informed by our understanding of current smart homes, we surface the social implications of current smart homes around access control and privacy in a home, the usability challenges of automation, and the tension between current approaches to address interoperability and reliability in smart homes.

4.1 Access Control

If access control in a smart home is not designed carefully, it could give one occupant more control and power over the devices in the home, compared to other occupants, which could affect interpersonal relationship in the home [24]. This raises an important design question: Should all occupants have equal access to the smart home? In non-smart homes, Johnson and Stajano argue that access policies for devices generally follow what they call the ‘Big Stick’ principle, which states whoever has physical access to a device is allowed to control that device [10]. Extending this principle to smart homes implies that an occupant should have at least the *operate* level of access for the devices in the same room. But this principle may not always apply, e.g., parents may not want their kids to access the smart TV during study time [18].

Tension: Current approaches to access control. The smart homes we tested used different access control models. SmartThings gave equal access privileges to all users, whereas Apple Home and Wink allowed for device-based access (i.e., primary user can grant other users access to certain devices). Vera provided more of a role-based access model, where the primary user defines roles for all other users, and access for each role is set by Vera. Between these two approaches (letting users decide for themselves vs. system enforcing equal access) it is not clear which one would be a better approach for a home setting. The equal-access-for-all approach (used by SmartThings) is simple and easy-to-understand for users, but does not support access control use cases like parental control or guest access, whereas the other approach (used by Apple Home, Vera, and Wink) supports some access control use cases but gives more power to the primary user.

Recommendation: Location-based access policy for guests. When a guest leaves, if their access is not revoked promptly, they can remotely access the devices in the home. The smart home owner may forget to promptly revoke a guest’s access or may hesitate to do so, for fear that it might appear rude to remove access as soon as the guest leaves the home. If an owner does remember to revoke a guest’s access, when the guest visits again, the guest would need to be granted access again. A location-based access policy, similar to what Apple Home provides, that leverages guests’ mobile phone to determine their location could grant guests access to devices only when they are present in the home. Such context-based access control policy could simplify access control for guest users, by automatically limiting their access only when they are at home.

4.2 Privacy

The privacy implications of a smart home are not limited to the individuals who are registered users of the smart home but also extend to *all* the smart home occupants as well as *guests* that visit the home.

Recommendation: Reduce privacy leaks via side-channels.

In a smart home, data is shared between users, and depending on how it is shared, privacy leaks may happen. For instance, consider device activity logs, which can help users understand their own behavioral patterns. If, however, the device activity log reveals device usage pattern of a particular user, that user could feel being watched. For example, in a shared smart home apartment, roommates may prefer if their trips to bathroom were not logged. Such privacy risks can be reduced by anonymizing device activity logs, storing activity logs only for a short period of time, or maybe for some devices the smart home could provide a “do not log” option.

Another example of privacy leak via side-channels is the misuse of the user-presence feature in SmartThings. In our experiments, SmartThings shared users’ presence status (determined by the location of the user’s phone) in the home with other users. Some may find location sharing convenient to let their family know when they arrived home, or to create automations based on their location; others may see it as an invasion of privacy. There is a need to design a usable permission model for location sharing that gives users more control over who can access and use their location.

Tension: Guest privacy. We usually think about privacy of home occupants from visiting guests, but what about the privacy of guests from smart home occupants? Consider guests who are staying for a few days (or longer) while the smart home occupants are away on a vacation; or consider an Airbnb host who rents his/her smart home to visitors. In this case, should the smart home owner have remote access to the smart home while it is occupied by a guest? And, how should the smart home owner retain ownership privileges but without invading the guest’s privacy? Disabling remote access to the smart home disables owner’s remote access, but also disables the guest’s remote access, which the guest may not want. So selectively disabling remote access for certain users may be desirable for such situations.

Tension: Utility vs. privacy with continuous sensing. Some smart devices, with their continuous sensing, can provide a lot of utility to users, but depending on where the data is stored, who has access to the data, and what that data is used for, there can be serious privacy implications for smart home occupants. For example, smart meters can provide users their detailed energy consumption and grid companies use this information to ultimately improve their service to the users, but this data can also be used to infer information about the appliances in the home and their use, and also users’ activities in general. Smart assistants like the Amazon Echo or Google Home are always listening to readily respond to users’ commands, and companies may store all the recorded audio to improve their service, but this data can also be used to infer information such as identifying the number of people in a room, guessing a user’s mood (happy, stressed, depressed), or building a model of the user’s speech; users of smart assistants may not be aware of these implications. And similarly, smart cameras (e.g., baby monitors, security cameras) record video and provide users with relevant notifications (e.g., baby woke up, a trespasser caught on camera), but companies may store the video, which can have serious privacy implications for the people captured in the videos. Some prior approaches to this problem span high-level privacy

guidelines for building ubiquitous systems [11], privacy-aware architectures for building smart home apps [8], and privacy-aware analysis algorithms [13], but there is need for effective solutions to reduce this tension and for mechanisms that enable users to make informed trade-offs.

4.3 Automations and Interoperability

When users create automations, they may not necessarily be aware of failure cases, particularly failure of the trigger device or the actuator device.

Recommendation: Ask for post-failure actions. When a user creates an automation, asking the user what the smart home should do when the automation fails may serve two goals: *i*) educating the user by informing that the automation could fail – something the user may not be aware of – and the system could offer some information on how that automation could fail; and *ii*) increasing user confidence in the smart home’s reliability, by educating the user and giving her more control [23].

Recommendation: Detect possible failures. To detect failures in an automation, a smart home should know the status of the input trigger device and of the output actuator device. In our tests, all the smart homes were aware of the status of all their devices. So, current smart homes could easily detect a device that becomes unavailable and identify the affected automations. However, identifying all the affected automations may not be trivial if there are automations that dependent on other automations. For example, if an automation, say Q, is dependent on the successful execution of another automation P, but if P’s trigger device becomes unavailable and P fails, should the smart home also mark Q as a failed automations and run Q’s post-failure action? The correct choice is not obvious. Thus, implementing failure detection and notification would require careful design.

Tension: Interoperability at the expense of reliability. Interoperability between smart home devices and smart home hubs is one of the main challenge for smart homes [6]. To deal with this challenge, some device vendors are choosing the cloud approach, where a device and the hub communicate through the cloud even when the device and the hub are on the same local network and could communicate directly (locally). This cloud approach allows vendors to make their devices compatible with different smart home hubs, available now and in near future, but this same approach also increases network dependency and introduces additional privacy concerns, as the smart home data may reside in multiple servers. Thus, although interoperability can be addressed by integrating in the cloud, this approach introduces new privacy risks and makes smart homes less resilient to network failures.

5 RELATED WORK

Much of the prior work on modern smart home devices has been focused on individual IoT devices or individual smart homes. For example, Yang et al. conducted a user-experience study with people living with Nest smart thermostat, and found that participants had difficulty understanding how the system worked, which lead to reduced interaction with the thermostat [23]. In 2013, Ur et al. conducted access control cases studies of three smart home products [21], and one of their devices was Philips Hue, which we

also used in our evaluation and we found similar access control issues with Philips Hue. Woo and Lim conducted a 3-week study with participants using DIY smart homes and they identified six stages of DIY-usage [22]; their stages are similar to the phases we used in our evaluation. And more recently, Fernandes et al. analyzed security of SmartApps in emerging smart homes [7], and Alrawi et al. analyzed the security of various IoT devices [2].

In contrast with these previous works with particular smart home products, we take a broader approach and study seven different smart homes. Through such a collective analysis we can learn the similarities and differences in the design choices that smart home vendors make, we can identify past research recommendations that are being used (or not being used) in current smart home products, which can help guide future research.

6 CONCLUSION

In smart home computing, the mobility of the users, the changing context and needs, and the continuous sensing in the home, raise several design challenges. In this paper, we systematically studied seven popular, commercially available consumer smart homes, to compare their design choices around access control, privacy, and automation, and to understand how these smart homes handle certain edge cases (e.g., broken automations). We found, for example, that smart homes are converging on design choices related to smart home's reliability (e.g., local processing and local control), but that their approaches for access control and privacy are different; that access control and data sharing policies in some smart homes could enable occupants to spy on other occupants; and that alternate modes of interaction (e.g., voice-controlled devices) add convenience but could undermine access control policies in the smart home. From our evaluation of these design points and failure cases, we surface key issues around access control and privacy, the usability challenges of automation, and tensions between interoperability and reliability. These lessons and design alternatives can help inform future research and next-generation smart home technologies.

ACKNOWLEDGEMENTS

We thank our shepherd, Aakanksha Chowdhery, as well as our anonymous reviewers, for their valuable feedback. We also thank Anna Kornfeld Simpson for helpful feedback on earlier drafts. This work was supported in part by the National Science Foundation under Award CNS-1565252, the MacArthur Foundation, and the University of Washington Tech Policy Lab.

REFERENCES

- [1] Gregory D Abowd and Elizabeth D Mynatt. 2000. Charting past, present, and future research in ubiquitous computing. *Proceedings of the ACM Transactions on Computer-Human Interaction (TOCHI)* 7, 1 (March 2000). <https://doi.org/10.1145/344949.344988>
- [2] Omar Alrawi, Chaz Lever, Manos Antonakakis, and Fabian Monrose. 2019. SoK: Security Evaluation of Home-Based IoT Deployments. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*. <https://doi.org/10.1109/SP.2019.00013>
- [3] A J Brush, B Lee, R Mahajan, and S Agarwal. 2011. Home automation in the wild: Challenges and opportunities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. <https://doi.org/10.1145/1978942.1979249>
- [4] A. J. Bernheim Brush and Kori M. Inkpen. 2007. Yours, Mine and Ours? Sharing and Use of Technology in Domestic Environments. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*. Springer-Verlag, 18. <http://dl.acm.org/citation.cfm?id=1771592.1771599>
- [5] Scott Davidoff, Min Kyung Lee, Charles Yiu, John Zimmerman, and Anind K. Dey. 2006. Principles of Smart Home Control. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*. Springer-Verlag, 16. https://doi.org/10.1007/11853565_2
- [6] W. Keith Edwards and Rebecca E. Grinter. 2001. At Home with Ubiquitous Computing: Seven Challenges. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*. Springer-Verlag, 17. https://doi.org/10.1007/3-540-45427-6_22
- [7] Earlene Fernandes, Jaeyeon Jung, and Atul Prakash. 2016. Security Analysis of Emerging Smart Home Applications. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*. <https://doi.org/10.1109/SP.2016.44>
- [8] Jason I. Hong and James A. Landay. 2004. An Architecture for Privacy-sensitive Ubiquitous Computing. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, 13. <https://doi.org/10.1145/990064.990087>
- [9] Iulia Ion, Niharika Sachdeva, Ponnuram Kumaraguru, and Srdjan Čapkun. 2011. Home is Safer Than the Cloud!: Privacy Concerns for Consumer Cloud Storage. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*. ACM, Article 13, 20 pages. <https://doi.org/10.1145/2078827.2078845>
- [10] Matthew Johnson and Frank Stajano. 2009. Usability of Security Management: Defining the Permissions of Guests. In *Trust, Privacy and Security in Digital Business*. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-04904-0_36
- [11] Marc Langheinrich. 2001. Privacy by Design: Principles of Privacy-Aware Ubiquitous Systems. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*. https://doi.org/10.1007/3-540-45427-6_3
- [12] Michelle L. Mazurek, J. P. Arsenault, Joanna Bresee, Nitin Gupta, Iulia Ion, Christina Johns, Daniel Lee, Yuan Liang, Jenny Olsen, Brandon Salmon, Richard Shay, Kami Vaniea, Lujo Bauer, Lorrie Faith Cranor, Gregory R. Ganger, and Michael K. Reiter. 2010. Access Control for Home Data Sharing: Attitudes, Needs and Practices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, 10. <https://doi.org/10.1145/1753326.1753421>
- [13] Stephen McLaughlin, Patrick McDaniel, and William Aiello. 2011. Protecting Consumer Privacy from Electric Load Monitoring. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM. <https://doi.org/10.1145/2046707.2046720>
- [14] Andrew R. McNeill, Lynne Coventry, Jake Pywell, and Pam Briggs. 2017. Privacy Considerations when Designing Social Network Systems to Support Successful Ageing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, 13. <https://doi.org/10.1145/3025453.3025861>
- [15] Sarah Mennicken and Elaine M Huang. 2012. Hacking the natural habitat: An in-the-wild study of smart homes, their development, and the people who live in them. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*. https://doi.org/10.1007/978-3-642-31205-2_0
- [16] Sarah Mennicken, Jo Vermeulen, and Elaine M. Huang. 2014. From Today's Augmented Houses to Tomorrow's Smart Homes: New Directions for Home Automation Research. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. ACM, 11. <https://doi.org/10.1145/2632048.2636076>
- [17] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman M. Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*.
- [18] Stuart Schechter. 2013. The User IS the Enemy, and (S)he Keeps Reaching for that Bright Shiny Power Button!. In *Proceedings of the Workshop on Home Usable Privacy and Security (HUPS)*.
- [19] D. K. Smetters and Nathan Good. 2009. How Users Use Access Control. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*. ACM, Article 15, 12 pages. <https://doi.org/10.1145/1572532.1572552>
- [20] Nick Statt. 2017. How an army of vulnerable gadgets took down the web today. Online at theverge.com. <https://www.theverge.com/2016/10/21/13362354/dyndns-ddos-attack-cause-outage-status-explained>
- [21] Blase Ur, Jaeyeon Jung, and Stuart Schechter. 2013. The Current State of Access Control for Smart Devices in Homes. In *Proceedings of the Workshop on Home Usable Privacy and Security (HUPS)*. <https://www.microsoft.com/en-us/research/publication/the-current-state-of-access-control-for-smart-devices-in-homes/>
- [22] Jong-bum Woo and Youn-kyung Lim. 2015. User experience in Do-It-Yourself-style smart homes. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. <https://doi.org/10.1145/2750858.2806063>
- [23] Rayoung Yang and Mark W. Newman. 2013. Learning from a Learning Thermostat: Lessons for Intelligent Systems for the Home. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. ACM, 10. <https://doi.org/10.1145/2493432.2493489>
- [24] Eric Zheng, Shrirang Mare, and Franziska Roesner. 2017. End User Security and Privacy Concerns with Smart Homes. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/zeng>