

Understanding and Improving Security and Privacy in Multi-User Smart Homes: A Design Exploration and In-Home User Study

Eric Zeng
ericzeng@cs.washington.edu

Franziska Roesner
franzi@cs.washington.edu

*Paul G. Allen School of Computer Science & Engineering
University of Washington*

Abstract

Smart homes face unique security, privacy, and usability challenges because they are multi-user, multi-device systems that affect the physical environment of all inhabitants of the home. Current smart home technology is often not well designed for multiple users, sometimes lacking basic access control and other affordances for making the system intelligible and accessible for all users. While prior work has shed light on the problems and needs of smart home users, it is not obvious how to design and build solutions. Such questions have certainly not been answered for challenging adversarial situations (e.g., domestic abuse), but we observe that they have not even been answered for tensions in otherwise functional, non-adversarial households. In this work, we explore user behaviors, needs, and possible solutions to multi-user security and privacy issues in generally non-adversarial smart homes. Based on design principles grounded in prior work, we built a prototype smart home app that includes concrete features such as location-based access controls, supervisory access controls, and activity notifications, and we tested our prototype through a month-long in-home user study with seven households. From the results of the user study, we re-evaluate our initial design principles, we surface user feedback on security and privacy features, and we identify challenges and recommendations for smart home designers and researchers.

1 Introduction

Smart devices and smart home platforms, such as Samsung SmartThings, Philips Hue lights, Google Home, the Amazon Echo, and Nest thermostats and cameras, are being increasingly adopted and deployed in the homes of end users. These devices and platforms allow users to remotely control and monitor their devices as well as to create automations (e.g., automatically locking the door when the user leaves home).

Security and Privacy in Multi-User Smart Homes. Smart homes are fundamentally multi-user platforms. Multiple people living in or accessing a home—including partners,

roommates, parents and children, guests, and household employees—may want or need the ability to use and configure the smart devices within the home. As prior work (e.g., [14, 36, 41]) has begun to show, conflicts and tensions may arise between these multiple stakeholders—even in generally non-adversarial (e.g., non-abusive) households. For example, the more tech-savvy users who install smart devices in their homes may intentionally or unintentionally restrict other users from accessing home functions (like thermostats) that were previously physically accessible [14, 41]; privacy concerns and violations may arise between co-occupants [3, 41]; and remote control of devices can be used for harassment [3].

Unfortunately, current smart homes are not yet thoughtfully designed for interactions between and use by multiple people. Though prior work has surfaced the need for additional access control options [16], transparency, and privacy features [41], many commercial smart home platforms present only simple security and privacy controls, or even none at all [22]. For example, Samsung SmartThings, a popular smart home platforms, forces home administrators to choose between provisioning additional accounts with administrator privileges or not provisioning additional accounts at all [33].

Designing to Address These Challenges. Providing multi-user smart home security, privacy, and usability is not a straightforward matter of simply building it, but rather requires careful consideration of a complex design space. We take a step back to ask: What security, privacy, and other goals *should* a multi-user smart home design aim to achieve? How might it achieve those goals? And do those goals and their implementation meet the needs of end users in practice?

In this work, we focus specifically on answering these questions for generally functional households without explicitly adversarial relationships. That is, we consider “typical” tensions that may arise between roommates, partners, and parents and children as they interact with and through a smart home; we do not consider explicitly adversarial relationships, such as domestic abuse. Addressing such challenging situations is of course also critically important, but we observe that even the seemingly “easy” case has not yet been sufficiently addressed

in prior work or today’s commercial platforms.

To begin answering these questions, we thus systematized prior work to develop an initial set of design principles for smart homes in generally non-adversarial multi-user households: *access control flexibility, user agency, respect among users, and transparency of smart home behaviors*. Based on these initial principles, we designed and prototyped a mobile app for smart home control, which includes concrete features such as location-based access controls, supervisory access controls, and activity notifications.

In-Home User Study and Design Recommendations. To evaluate how well our proposed design principles and our prototype’s specific design choices meet the use cases of real (non-adversarial) households — and to gain a deeper understanding of the multi-user smart home access control needs and use cases of this class of users — we conducted a month-long *in situ* user study. We deployed our prototype with seven households in the Seattle metropolitan area.

The empirical findings from our user study allow us to evaluate and refine our proposed principles for security and privacy in multi-user smart homes, and we surface technical directions and open questions for platform designers and researchers, which were not apparent in prior work that did not conduct *in situ* design evaluations.

Among our multiple findings, we found that for some of our participants, positive household social norms and relationship dynamics obviated the need for technical access controls. This finding suggests directions and questions for future work, including: How can a smart home platform design leverage or scaffold these social norms rather than simply existing alongside them? And how can the platform simultaneously support use cases and user groups where these social norms and relationship dynamics are not as positive [3] or (as in the case of our participant families with teenagers) in tension?

Another finding surfaced through our user study was that participants’ varied access control desires required our prototype to support complex combinations of access control options. Unfortunately, when we increased complexity, it decreased usability, potentially discouraging less motivated or savvy users from using access controls. This finding raises the question: how can smart home designers increase the flexibility of smart home access control systems while making the complexity manageable for all users?

Contributions. Our work makes the following contributions:

1. *Design Principles and Prototype:* We systematize from prior work a set of possible design principles for security and privacy in multi-user smart homes, and we develop a prototype based on these principles targeting generally cooperative households.
2. *In-Home User Study:* We use our prototype to conduct a month-long in-home user study with seven (non-adversarial) households, including couples, roommates, and families with children of various ages. Our study

serves to both test our proposed design principles in practice and to more generally enrich the literature on people’s security and privacy needs, concerns, and priorities in a multi-user smart home.

3. *Lessons and Recommendations:* Based on our design experience and in-home study, we reflect on our proposed design goals for multi-user smart homes and surface future technical directions as well as open questions for designers and researchers.

2 Background and Motivation

Smart homes raise significant potential security and privacy challenges. These challenges include, for instance, vulnerabilities in the devices themselves (e.g., [29]) and privacy concerns due to ubiquitous recording in the home [6, 21, 42].

In this work, we focus primarily on multi-user security and privacy: how peoples’ behavior and usage of the smart home can impact each others’ security and privacy. We begin by systematizing the multi-user security and privacy issues prior work has identified for smart homes, as well as the shortcomings of existing approaches in addressing these issues.

2.1 Multi-User Challenges in Smart Homes

Prior work suggests that smart homes can cause or intensify conflicts or tensions between people living in the home — even when relationships between people are not explicitly adversarial (e.g., abusive).

Power and Access Imbalances. One negative dynamic that emerges from smart homes is a power imbalance between the person(s) who install(s) and configure(s) the home, and the other users who are more passively involved. In the worst-case — in the context of intimate partner violence — abusers may have total control over the smart home, enabling harassment and abuse [3]. However, power imbalances also arise in more benign relationships. For example, Geeng et al. observed how more tech-savvy users have more agency in the home, including more access to device functionality, more information about what devices and people in the home are doing, and the power to restrict others from using devices [14].

Privacy Violations. Smart homes can also intentionally or unintentionally used to expose privacy sensitive information about one user to another. Zeng et al. found examples of such situations, like users being unaware of automated notifications sent by cameras to their landlord, and users feeling a loss of privacy because others could view their behavior through smart home logs [41]. Choe et al. studied how devices that capture video, audio, and other behavioral traces could cause tensions between household members, or between guests and household members, who would object to being recorded [6].

Direct Conflict. Lastly, smart homes can be focal points of conflict between people in the home, both due to explicit

malice (e.g., abuse) and due to ordinary conflicts between household members. For example, prior work has documented conflicts arising due to differences in opinion on thermostat setting [14, 41], due to conflicting goals between parents and teens in the context of entryway surveillance [36], or due to the potential use of recorded evidence in household disputes [6].

2.2 Additional Actors: Apps and Automations

The above multi-user issues are compounded by the presence of additional “actors” in smart home systems: third-party apps and integrations that users may install (such as SmartApps or IFTTT), as well as end-user programmed automations. These apps and automations can range from simple rules (such as automatically locking the door or turning off lights when leaving home) to more complex “smart” features that integrate with other cloud services, e.g. weather data and calendars.

These applications and automations can expose users to physical security risks and privacy violations. Third-party applications and automations may be expressly malicious, or buggy and exploitable (e.g., [11]). Moreover, end users themselves may make mistakes programming automations, leading to unexpected behavior, bugs, and potential security and privacy risks [27, 34, 39]. In a multi-user smart home, this combination of actors means that when something unexpected happens in the home, it may be challenging or impossible for a user to determine whether it was the result of another user actuating the smart home remotely, a buggy application or automation, a legitimate application or automation that another user installed, or explicitly malicious activity.

2.3 Shortcomings of Existing Approaches

Though many commercial smart home platforms exist, and a growing body of research literature supports the need to address the above challenges, we are not aware of existing approaches that succeed at addressing them and/or have been rigorously evaluated with end user — neither for explicitly adversarial settings nor in generally cooperative households.

There are many types of access control policies that could be used in smart homes, including time-based policies, location-based policies, per-user policies, and per-device policies. However, Mare et al. found that adoption of these techniques in smart home platforms is uneven and limited [22]. Some platforms support a subset of these policies, e.g., Apple Homekit has location-based access controls, and Vera has multiple privilege tiers for admins and guests. However, some popular platforms have minimal or no access control at all: Samsung SmartThings has only a single privilege level for all users and no access control policies, while Google Home and Amazon Echo do not authenticate voice commands. He et al. [16] and Ur et al. [35] found similar fragmentation of access control and authentication policies between individual devices: some devices like door locks had many access

controls, while others like smart thermostats had none.

While having no access controls or user roles at all is clearly insufficient for user needs (e.g., [33]), the jury is still out on what are the *right* access control designs for multi-user smart homes. To that end, He et al. [16] surveyed hundreds of participants to understand their smart home access control preferences, such as which device capabilities people felt need restrictions (like “deleting door lock logs”) and which types of device capabilities and people could use special contextual controls (e.g., allowing children to control devices only when parents are around). These survey results provide a valuable basis for future smart home access control designs, but they still only represent a theoretical view of people’s preferences. To the best of our knowledge, there have been no direct, *in situ* evaluations of multi-user smart home access controls designs with end users. We aim to close that gap in this work.

3 Scope and Research Questions

Prior work has surfaced many multi-user security and privacy challenges in current smart home systems. However, this body of research lacks concrete design proposals that have been evaluated with end users. We aim to advance our understanding of this space.

We focus in this work on generally functional multi-user households, rather than on explicitly adversarial situations (e.g., domestic abuse) or cases where users do not belong to a household together (e.g., Airbnb-style rentals). Understanding and designing for these cases is also critical, but different (and significant) challenges exist in designing systems that are resilient to motivated adversaries with malicious intent, elevated privileges, and physical device access [23]. We discuss the ways in which our work may address — but also falls short of addressing these challenges, in Section 7.4.

Yet prior work has not answered the question of how to design multi-user smart homes for “typical” households; thus, in this work, we seek to answer two research questions:

RQ1: How should a smart home be designed to address multi-user security and privacy challenges (in generally functional households)? What design principles and concrete features may help mitigate tensions and disagreements among otherwise cooperative (e.g., non-abusive) co-habitants that stem from multi-user security and privacy issues?

RQ2: What security and privacy behaviors and needs do these smart home users exhibit in practice? Prior work has provided some understanding of users’ security and privacy preferences in the smart home, like preferences for access controls [16], or examples of undesirable situations [14, 41]. However, these preferences could conflict with other priorities, such as utility and convenience. We ask: when presented with a smart home with more advanced security and privacy features, how do people (in non-adversarial households) use them in practice? Do users’ security and privacy related be-

aviors differ from their stated preferences? Do our initial design principles match their needs?

To answer these research questions, we take a two-part approach. First, we design and implement a multi-user smart home interface, based on design principles (Section 4.1) that we distill from prior work. Second, we conduct an in-home user study using our prototype, to evaluate whether these design principles meet user needs in practice, and to improve our understanding of users’ behaviors given improved multi-user security and privacy features in a smart home.

4 Prototype Design and Implementation

To support the investigation of our research questions, we prototyped a mobile application for controlling smart homes that provides multi-user security and privacy features such as access controls, designed for households in which members are generally motivated to cooperate. We now describe the guiding design principles for our prototype.

4.1 Initial Design Principles

We developed our prototype based on lessons from prior work, which suggested that the following design principles may be important for multi-user smart homes:

Access Control Flexibility. Prior work [16] has suggested that smart home access control and authentication systems should be flexible enough to support a wide variety of use cases, people, and types of relationships that exist in homes. We aimed to support a variety of relationships, like couples, roommates, children, guests, and domestic workers, and also different contextual factors, like location. These factors can be combined to create the policy that suits the user.

User Agency. Prior work [14] found power imbalances among smart home users that reduce the agency of users with less (technical or interpersonal) power. We aimed to support a feeling of agency for all users in the smart home, by making the smart home more accessible and discoverable. For example, for access controls, our prototype allows people to “ask for permission”, rather than to be locked out entirely. We aim to make smart home functionality more discoverable, by showing users which devices are nearby and accessible. We also aimed to simplify the process of on-boarding new users.

Respect Among Users. Prior work has surfaced significant potential for tensions and conflicts among users of a smart home (e.g., [14,41]). We aimed to encourage respectful usage of the smart home by minimizing conflict points: for example, making it harder for one user to remotely control or automate devices in a way that would surprise or disturb another.

Transparency of Smart Home Behaviors. Prior work suggests that smart home automations and apps may malfunction or act maliciously (e.g., [11,34]), violate the privacy of

unaware users (e.g., [41]), or confuse users who did not configure them. When smart homes are used for domestic abuse, abusers have harassed victims with remote control, masking it as automatic behavior [3]. We aimed for the smart home to transparently surface its behavior to all people in the home (realizing that there may be privacy implications, as we discuss below), especially when people are remotely controlling it, or when an automation/third-party app is acting on its own.

4.2 General Design Description

We designed a mobile application that allows multiple users to control their smart home devices. In terms of threat model, we assume that the control application and the underlying smart home (SmartThings, in our study) are trustworthy and uncompromised. We assume that third-party smart home automations or applications may be buggy or compromised, but our design does not aim to prevent such issues. We assume that users may use or configure the smart home in ways that are undesirable to others in the home, though we focus on cases in which this behavior is accidental or mildly malicious (e.g., “trolling”); we do not attempt to defend against a determined, malicious adversary (e.g., an abuser).

The basic interface of our app is similar to other mobile apps for controlling smart homes (e.g., Samsung SmartThings). The main view of the app displays a list of devices and their current status (Figure 1a). Devices can be organized by room for convenience. The state of a device can be adjusted by tapping its status, and tapping its name reveals options for access controls and notifications (described below).

We aimed to simplify the process of onboarding additional users, towards meeting the “user agency” principle. The first user must create an account with a username and password, but they can add other users by scanning a QR code on the new user’s phone. These additional users do not need a login, instead using public key authentication tied to their device.

4.3 Access Controls

Towards meeting the “access control flexibility” and “respect among users” principles, we designed access controls for accessing device capabilities, based on access control preferences and use cases surfaced in prior work (e.g., [10,16]).

Role-Based Access Control. Each household member has a separate user account. Users can be restricted from using a device via the ‘Allowed Users’ setting (Figure 1b). Users are also assigned to roles (admin, child, guest). Only admins have the ability to make configuration changes: changing access control policies, adding new users, organizing the devices.

Location-Based Access Control. Users can also be restricted from controlling device capabilities if they are not physically near the device, or not at home, using the “Remote Control” permission (Figure 1c). This access control can be

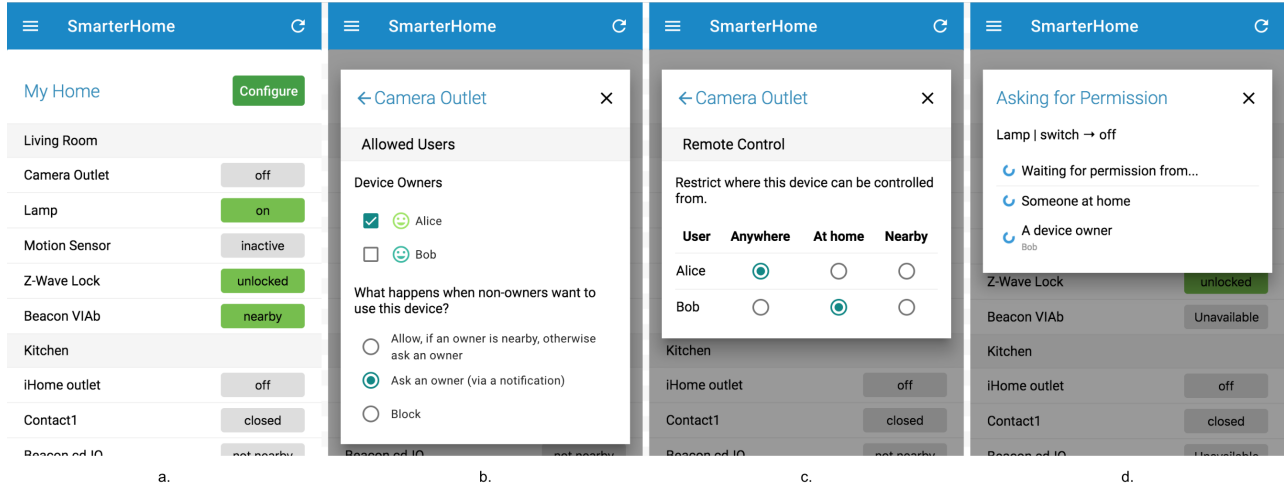


Figure 1: **Access Control UI.** From left to right: (a) The main interface for controlling devices. (b) Interface for setting access controls on devices, by role, and options for reactive/supervisory access control. (c) Interface for setting location-based access controls on a device, for each user. (d) Reactive access control prompt: what users see while waiting for approval.

set per-user, to accommodate use cases like only allowing guests and domestic workers to access smart home devices while in the house. It could also be used to promote respect among users by preventing them from remotely controlling devices like lights when other people are in the room.

Supervisory Access Control. Access controls are in some ways antithetical to user agency. For example, parents may want to use parental controls to keep children from causing trouble, but may not want to block children from using the smart home at all times, like when the parents are at home and are able to supervise. To serve this potential use case, we implemented supervisory access control (first proposed by He et al. [16]): if a user is restricted from controlling a device, they can still be permitted to control it if another (authorized) user is nearby (Figure 1b).

Reactive Access Control. Access control policies based on role and location could be too rigid for every situation. There may be occasional edge-cases where it does not make sense to enforce a policy. Towards the principles of increasing flexibility and supporting user agency for restricted users, we implemented reactive access control [10, 24]. If a user attempts to access a capability they do not have permission to use (Figure 1d), the app will ask a more privileged user for permission in real-time, by sending a notification to asking them to approve or deny the request (Figure 2c).

4.4 Activity Notifications

Towards meeting the “transparency of smart home behaviors” principle, i.e., to make it more transparent when the smart home is being remotely controlled, or controlled by automations and apps, we designed notifications that alert users when the states of home devices change. Each notification displays

the name of the device, the change in state, and the user or process responsible for causing the change (Figure 2).

We chose to use notifications over other designs that focused on visualizing automations and events in-app [5, 26], to explore a different point in the design space. Rather than having users navigate to a particular interface when motivated to investigate activity in their smart home, we hypothesized that real-time notifications could provide information in a more timely and relevant manner.

Because the number of notifications from the smart home could be overwhelming, we allowed users to disable notifications on a per-device basis, or to only receive notifications from physically nearby devices.

4.5 Discovery Notifications

Prior work (e.g., [14, 41]) suggests that one challenge with multi-user smart homes is that less technically savvy or engaged users may struggle with accessing smart devices. Thus, towards meeting the “user agency” principle, we wanted to make it clear which smart devices were nearby and could be actuated, especially for novice users. We designed a persistent notification which displays the status of nearby devices, and includes action buttons to toggle those devices (Figure 2b). This design makes devices that are nearby (and potentially relevant) accessible without needing to open the app. We designed it to be minimally intrusive—the notification is silent and is minimized at the bottom of the notification tray.

4.6 Implementation

We implemented a prototype mobile app with these features for Android, iOS, and web, using the Cordova framework. Rather than implement our own smart home controller that




Figure 2: **Overview of Notification Types.**

(a) *Activity Notifications.* When an event occurs in the home, this notification shows the name of device, capability being changed, and who or what caused the change.

(b) *Discovery Notifications.* Persistent, low priority notification that shows nearby devices and their current state; can be expanded to reveal action buttons for controlling devices.

(c) *Reactive Access Control prompt.* Appears when another user asks for permission to use a restricted device capability.

interfaced with hardware devices directly, our prototype connected to devices via the Samsung SmartThings API. Participants set up their smart home devices using SmartThings, and then used our app to control the system. Our prototype did not support automations and third party apps — users accessed this functionality through the SmartThings app. Our prototype consisted of 10257 lines of JavaScript, CSS, and HTML.¹

Proximity Sensing. To enable room-scale proximity-based features (location-based access controls, proximity-scoped notifications), we incorporated Bluetooth Low Energy beacons into our system. Beacons broadcast an ID that can be scanned by modern smartphones that support Bluetooth 4.0+. Users register physical beacons in our app using an ID printed on the device, and then assign it to a room in the app. When a user’s phone detects the beacon, the app infers that the user is near the devices in that room. We chose beacons as our proximity sensing solution out of convenience: they are supported by all

¹The source code and a demo of the prototype are available at <https://github.com/UWCSESecurityLab/smarter-home>




Figure 3: **Prototype Architecture Diagram.** We use the SmartThings API to communicate with smart home devices.

modern Android and iOS devices. However, our design does not require a specific proximity sensing technology; others such as WiFi or ultrasonic sensing would work as well.

SmartThings and iOS Limitations. Due to the limitations of the SmartThings API, activity notifications cannot attribute changes in home state to particular third-party apps, automations, or manual actuation of devices. For state changes in these categories, our implementation only displays “Triggered by an automation or manually”. Discovery notifications were only implemented on Android, as the iOS notification center does not support persistent, low priority notifications.

5 User Study: Goals and Methodology

Our prototype allows us to study the research questions we set out in Section 3. To do so, we recruited seven households in the Seattle metropolitan area to use our prototype to interact with their smart homes for a month-long period. We conducted studies between October 2018 and January 2019.

User Study Goals. Our goals in conducting the user study were two-fold, corresponding to our two research questions. First, we aimed to evaluate how participants used and reacted to the specific multi-user smart home features (and corresponding design principles) we implemented in our prototype. Second and more generally, we aimed to understand the multi-user access control and other needs and behaviors of end users, grounded in the use of a specific prototype in real homes.

Our specific evaluation questions, paired with the design principles our prototype intended to embody, included:

1. *Access Control Use Cases:* Is our current combination of access controls sufficient for users’ desired access control use cases? If not, what use cases are we missing?
2. *User Agency and Respect:* We envisioned that location-based and reactive access controls could be used to mitigate conflicts and tensions over controlling the home. Can we observe this in practice?
3. *Transparency of Smart Home Behaviors:* We envisioned that notifications could improve users’ mental models of smart homes, which would help with understanding privacy implications; and also improve security by creating a simple mechanism for auditing automations and apps. Do notifications provide these benefits to users in

practice? Conversely, do notifications harm privacy by revealing one person’s activity to other people?

Study Overview. We conducted a month-long *in situ* user study in the homes of participants. We recruited households in the Seattle metropolitan area. We provided a Samsung SmartThings smart home to households that did not already own smart home devices, or integrated SmartThings with the smart homes of households that owned an existing system. We collected qualitative data about participants’ previous experiences with smart homes and feedback on our prototype through interviews, experience sampling, and log data.

Recruitment. We recruited seven households, containing 19 participants who actively participated. Participating households were recruited via Facebook ads, targeted at people interested in smart homes and home DIY projects. People who clicked on the ads filled out a short survey including information about their household composition and interest in smart homes. We did not require participants to own any smart home devices prior to the study. We conducted a screening call with participants that met our criteria to collect additional information. We selected participants who lived within a 45 minute radius from our homes (so that it was feasible to make an in-home visit), and we aimed for a variety of multi-person household compositions, including roommates, families, and couples. Participating households are summarized in Table 1.

A limitation of our recruitment strategy and study design is that it introduces self-selection bias: our participants were likely to be living in generally cooperative households, with one or more technology early adopters. We discuss this, and other limitations, further in Section 7.5.

Initial Interview. We made an initial visit to participants’ homes to conduct a semi-structured interview about their existing experiences and attitudes towards multi-user smart home security, privacy, and usability issues (see Appendix A).

Following the interview, we assisted with the setup of any devices if needed, and then we set up our prototype app. We guided them through app installation because it required using the developer mode in SmartThings, which was cumbersome and not representative of a typical install experience for commercial apps. We also assisted participants in adding other household members, to ensure that we could study multi-user interactions (rather evaluating the onboarding barrier).

We also walked through the access control and notification features of the app, and collected their initial impressions of the features. To counteract participant response bias [8] we stressed that we were testing an imperfect prototype, and that we wanted honest, negative feedback on things that were not useful or usable. We used some participant feedback from this stage to iterate on our implementation and push updated features to participants throughout the duration of the study.

Daily Usage. Participants then used the app for 3-4 weeks during their daily lives. During this period, the integrated experience sampling interface in our app prompted participants

to provide feedback or to share anecdotes about multi-user interactions in the home. We also collected log data about how users set up access controls, permissions, and notifications.

Exit Interview. At the end of the usage period, we conducted a phone interview with each household. In this semi-structured interview we collected specific feedback about their experience using (or not using) the access control and notification features in our prototype. We also followed up on any interesting data from experience sampling or logs. A list of interview questions is available in Appendix B.

Compensation. Participating households were compensated \$250 over the course of the study, in installments. Participants could keep the provided smart home devices after the study, or return them for the equivalent cash value.

Ethics. The study was approved by the University of Washington’s human subjects review board. Participants had to be age 18+ to consent to participating; household members under 18 could participate with verbal assent and approval from their parents and guardians. We had approval to collect incidental data via the smart home on children who declined to participate or were too young to actively participate.

During the study, we experienced a security breach due to a firewall and database misconfiguration, resulting in the possible exposure of hashed passwords, log data, and temporary access tokens. Based on access patterns, we believe the data was accessed by port scanners, and not by targeted attackers. We remediated the issue within 24 hours of discovery. We notified our institution’s human subjects board, and contacted participants with a description of the issue, protective steps like changing matching passwords on other sites, and the option to opt out of the study. No participants opted out.

Analysis. We transcribed and analyzed 633 minutes of content from the 14 initial and exit interviews. We analyzed the interviews using a collaborative qualitative coding technique. First, two researchers read over all of the data and developed a codebook, using descriptive codes like “access control: use cases”, “relationship: guests”, “notifications: too noisy”, and “access control: trust/respect” (see Appendix C for a full list). Two researchers independently coded two interviews, and then met to resolve differences and clarify ambiguities in the codebook. Then, one researcher coded the remaining interviews based on our revised understanding of the codes. We used a custom code aggregation tool to help identify patterns and extract higher level themes across interviews.

6 User Study: Results

We now present the findings from our user study, including direct feedback on the features implemented in our prototype, and general findings about participants’ desired features and use cases, surfaced by their concrete experiences with our prototype and the smart devices in their homes.

Participants	Gender	Age	Devices	Household Info	
H1	H1A H1B	25-34 35-44	F M	Lock*, motion sensors*, contact sensors, thermostat*, security camera*, lights*, Amazon Echo*	Family with two non-participating children (0-6), living in house
H2	H2A H2B	25-34 25-34	M F	Lights*, Amazon Echo*, contact sensor, door lock (not connected)*	Couple, living in house
H3	H3A H3B H3C	25-34 25-34 25-34	F F F	Lights, contact sensor, motion sensor, power outlet	Roommates, living in apartment
H4	H4A H3B	25-34 25-34	F M	Lights, contact sensor, power outlet, Amazon Echo*	Couple, living in apartment
H6	H6A H6B H6C	35-44 45-54 13-17	F F M	Lights, contact sensor, door lock, Amazon Echo*	Family with 2 children (one aged 7-12), living in house
H7	H7A H7B H7C	18-24 18-24 18-24	F F F	Lights, contact sensor, motion sensor, power outlet, Ring video doorbell, Amazon Echo*	Roommates, living in house
H8	H8A H8B H8C H8D	45-54 45-54 18-24 13-17	F M F M	Lights, contact sensors, security cameras*, Amazon Echo*	Family with 2 participating children, one non-participating child (7-12), one non-participating relative (13-17), living in house

Table 1: **Summary of Participating Households.** Some children were too young to actively participate in the study. Asterisks (*) indicate devices households owned prior to the study.

6.1 Desired Access Control Use Cases

We begin by exploring the situations where participants *wanted* multi-user access controls, and what form of access control mechanisms participants wanted. In some cases, our prototype was able to fulfill participants’ goals, and in others, the ability to explore concrete access control features in the context of their own home evoked hypothetical policies that they felt would better suit their needs.

Location Restrictions for Visitors. H1A wanted an access control setting that would allow visitors like guests and domestic workers to be able to access and control the devices in her home, but only while they were physically present.

I don’t want the nanny, who’s here all day — I trust her, obviously, or she wouldn’t be with my kids — but at the same time, like I don’t necessarily need her to be at her house, being able to control the lights at my house. ...if I have guests coming into my house, I’d like them to control automations, but... I certainly don’t want them having admin control. I’d prefer to have them to have geofenced control. (H1A-Initial)

At the time, our prototype’s location-based access controls did not quite meet her requirements, because it could only be applied as a blanket policy for all users of a given device. Based on this feedback, we updated the prototype to support location-based access controls both per-user and per-device.

Preventing Configuration Changes. Some participants were concerned about other family members accidentally

making changes to access control policies, automations, or device configuration. H1A recalled when they set up their smart home, H1B (her spouse) caused confusion by accidentally pairing some devices multiple times. As a result, H1A set H1B at the child privilege level in our prototype, which prevented him from configuring access controls and rooms.

H8A did not want her children to either change or override the existing automation for the porch light, which turned the lights on automatically at night for security purposes, nor did she want them to be able to change access control policies. As a result H8C/D were set at the child privilege level in our app (and were also not added to the native SmartThings app, from which the automations were created).

Parental Controls for Device Usage. Parents in our participant sample expressed interest in placing restrictions on children to prevent mischief or other undesired uses of devices. For example, H1A and H8A wanted to restrict their children from turning on/off security cameras. However, participants did not use our prototype’s features for restricting access to any devices in practice, for reasons we discuss below.

A parental control goal that we did not anticipate was that H1A and H8A were more interested in using the smart home to regulate screen time, e.g. blocking internet access at certain times, and using a smart power outlet to turn off the TV.

Devices in Private Rooms. The roommates of H3 placed smart light bulbs each of their bedrooms, and set an access control policy so that only the room’s owner could control the lights. They reported that it was “comforting” and a “good

feature to have” (H3C), but that in practice, they never encountered the access controls because they were respectful of each other and did not ever attempt to control another person’s lights. (We discuss similar cases of social norms obviating the use of technical access controls below.)

Preventing Remote Access for Media Devices. H4A/B expressed interest in location-based access controls for their Amazon Echo, based on past experiences where one of them accidentally changed the audio that was playing from outside of the room or house, due to confusion over whose Bluetooth device or Spotify account was playing. We did not see similar interest in location-based access controls for other device types — perhaps because unlike lights or locks, which are useful to remotely control for security and energy saving purposes, media devices are only useful to the people physically in the room.

Access Controls for Voice-Controlled Devices. H8A became aware that their Amazon Echo could be used to bypass the access controls and authentication of our prototype (see Section 7.3 for more detail). In one instance, she used this to allow her mother-in-law to access the smart home without installing our prototype. However, she also wanted the Echo to authenticate users by voice, so that they could use access control policies for to their youngest son, who was too young to have a phone but could control devices via the Echo.

6.2 Reasons for Not Using Access Controls

Based on findings about multi-user smart home tensions in prior works, we expected that households would use our access controls, for at least some of the potential use cases outlined in our design principles (Section 4.1). However, in general, we found that the access controls we implemented did not fit with the participants needs and use cases.

We analyzed participants’ usage logs, and found that while most households experimented with using access control policies in the first few days after the initial interview, most of them quickly settled on the least restrictive access control setting, not continuing to use location-based, role-based, or supervisory access controls to restrict access to devices. The only household that kept any access controls enabled was H3, a household of roommates who enabled per-device role-based access controls on the lights in their private bedrooms. However, none of the roommates ever attempted to violate these access controls (i.e., tried to turn on or off each others’ lights).

Given this limited long-term use of access control features in practice, we thus focus on our qualitative interview data, to dig into the reasons *why* participants did not use the access controls more than they did. Our findings surface several reasons that are more fundamental than simply reactions to our specific implementation — i.e., reasons that participants may not have used *any* access controls, regardless of design.

Social Norms, Trust, and Respect. The most common rea-

son participants cited for not setting access controls was trusting each other enough that they were not concerned about device misuse, relying instead on established household and interpersonal norms. We observed such trust and norms among relatively equal relationships, like partners and spouses:

No, we didn’t turn [remote control restrictions] on either... We both wanted full permissions to do anything whenever, we weren’t worried about the other. I had no concern that H2B, from not nearby, would turn off the lights. (H2A-Exit)

We also observed trust and norms among roommates:

I think we’re all pretty respectful and we wouldn’t turn on and off each others’ lights. (H7A-Exit)

And even with children:

If [H6C] were a different kid, I would probably leave [remote control] turned off for him. But for him, it would be useful, I would turn it on for him. ...I think it’s going to be very specific to who is using it, and having the option is important, but he’s just very responsible, so it could’ve been handy for him to be able to do something from school, like turn on and off lights. (H6A-Initial)

Participants mentioned similar social norms about multi-user privacy. For example, H1A and H8A/B were aware that it was possible to eavesdrop using devices like the Amazon Echo or security cameras, but chose not to do so.

Interference with Other Functionality. Particularly with location-based access controls, participants often felt that the available access controls were too restrictive and prevented them from accomplishing other goals. In our initial design, we expected that location-based access controls could serve a number of goals, like access control for guests, or preventing mischief or inconsiderate use of remote controls. However, multiple participants wanted unfettered remote control access, particularly for lights, because it was convenient.

I think a big thing for us was in case we forgot to turn off the lights or something, that was like the appeal, to turn it off remotely. (H7A-Exit)

Like the times when we would both need access to turn off the light we forgot to turn on, were more frequent than any need to restrict us from being able to remotely control it. (H4B-Exit)

In other words, at least for the smart devices our participants had, the convenience for all members of the household to be able to exercise remote control outweighed any concerns about intentional or accidental remote misuse.

Lack of Concern About Devices. Participants did not feel concerned enough to use access controls for certain types of devices, or for devices in certain locations. For example, participants did not feel that smart lights were sensitive enough for access control (but did want restrictions on more sensitive devices, such as cameras, for guests and children).

We cannot say whether participants would have used more access controls for more sensitive devices — since we allowed our participants to select their own devices, their *a priori* threat models likely influenced their devices selection (i.e., selecting devices they were comfortable having in their home). We discuss this issue further in Section 7. Moreover, we note that these limited multi-user concerns were consistent with participants’ overall smart home related threat models (likely due to self-selection bias). Though some participants were aware of potential risks such as password compromise, vulnerabilities in wireless protocols, data collection by companies, or lost phones, they did not consider these risks to overwhelm the utility of the smart home.

Some participants also did not find it necessary to control access to devices located in household common spaces, like locks and lights — again showing physical-world household social norms reflected in the configuration of the smart home.

6.3 Limited Utility from Activity Notifications

We found varied use of activity notifications among our participants. From our log data, we observe that 14 participants had activity notifications on at all times for all devices, while 4 participants used a combination of settings: on, off, and proximity scoped for various devices. This data suggests that proximity scoping provided utility for some participants. (One child participant did not have the app installed.)

But having notifications enabled does not necessarily mean that participants found them useful; we now dig further into our qualitative interview data to understand whether and how the notifications were useful to participants. Our participants found notifications useful for a few specific use cases, like home security and sanity checking their smart home automations. However, we did not find much evidence that our notifications provided benefits for transparency and agency.

Monitoring and Home Security. Participants found notifications to be most useful for home security and monitoring purposes. H1, H6, H7, and H8 used our prototype’s notifications in conjunction with sensors on their exterior doors and windows, to passively monitor their home’s security. H3C used notifications to monitor devices in their bedroom, to check if others were entering the room.

Proximity Scoping for Activity Notifications. While participant H8D found proximity scoping useful, as she did not want to be notified about devices while away from home, other participants said that the feature would be more useful if they could be notified only when *not* at home — either as a home security measure (H8A), or because they could already tell when their devices changed while at home (H4B).

Confirmation of Home Behaviors. Some participants found the notifications to be comforting because they confirmed that both people and automations were behaving as expected.

It was nice to know it was at that point in the day,

and really what I had it set on were essentially the lights to come on and go off at appropriate times, and so it was a notice that, yes, today is progressing as it should. (H6A-Exit)

Desire for Contextual Notifications. Our activity notifications prompted participants to propose more advanced, context-dependent notifications that would be more useful to them. For example, H3C suggested notifications which would suggest turning off the lights to save energy. H6A wanted more intrusive notifications when something incorrect happens (e.g., a window is open when it should not be).

We were not able to test whether notifications would be helpful for identifying actions caused by specific automations, because limitations of the SmartThings API did not let us see *which* automation caused an event to happen. None of our participants mentioned encountering a situation in which they wanted more specific information about provenance.

Quick Access via Discovery Notifications. Most participants did not notice or see discovery notifications. (Unfortunately, persistent notifications of this sort are not supported on iOS, and few of our participants were Android users.) One participant, H8D, was interested in these notifications, but for convenience, not device discoverability, as it allowed him to toggle the state of the device without opening the app.

Limited Concern about Privacy. No participants reported that the notifications affected their sense of privacy, nor that they changed their behavior as a result of knowing that notifications would be shown to others. Participants also did not report learning new information about others via notifications.

Notifications Were Overwhelming or Not Useful. For some participants, the notifications were annoying and overwhelming. H1A said she just did not care when other people, like her husband and nanny, used devices. H7A complained about redundant notification: each time someone walked through the front door, their doorbell and contact sensor would both trigger notifications, resulting in four notifications.

Other participants said that the notifications were not useful when at home, because it was information that was already apparent. Participants in H3 and H4 lived in small apartments, and could naturally observe all of the information from the notifications (e.g., the sound of others walking around and the glow of lights in other rooms). And H7A said that their dogs already notified them when people were at the front door.

6.4 Usability and Configuration Complexity

Hands-on experience with our app revealed that the complexity of access controls and other smart home features were adversely affecting the usability of the system. The complexity came from both the granularity of the settings, and the number of different devices managed by the home.

Complexity as a Barrier to Access Control Use. While we

aimed to make our prototype’s access controls as easy to understand as possible, the inherent complexity in the matrix of options may have still been too much of a barrier for novice users to configure them. For example, usability may have been an issue for H8A/B, where both expressed interest in setting various access controls during the feature walkthrough in the initial interview, but did they did not ending up using them. When we asked about other goals they might have for access controls and the smart home in general, H8A said:

It interests me, but you have to think it through, what you want to do, how it would benefit you... part of the Smart Things is you’re taking on a bit of a responsibility, getting it set up, getting it working, it’s kind of like getting a new computer, but there’s a bit of the downside, you have more options but it’s complicated. (H8A-Exit)

Design Complexity from Combinations of Settings. During the study, participants requested more fine-grained options for the access control and notification features. Based on this feedback, we iterated on the implementation of our prototype and released updates. However, we struggled with adding these features, as each additional access control dimension compounded the complexity of the interface.

One example was for location-based access controls. Initially, these access controls were set per-device. However, H1A and H8B wanted to set these access controls per-user in addition to per-device, so that they could restrict their nanny and kids (respectively), but not themselves. To fulfill this request, we had to surface more options ($3n$ options per device, where n is the number of users, instead of 3 options per device). As another example, if we wanted to add toggles for supervisory and reactive access controls to location-based access control when users are not nearby and try to use the device, there would not be enough space to display these options without an additional submenu, making it more laborious to set policies for each user and device (see Figure 1c).

Usability is fundamentally in tension with the desire to support access control flexibility and surface all of these options to users - we discuss this issue further in Section 7.

Displaying Access Control Policies. Participants remarked that it would have been helpful if the main device control page (Figure 1a) surfaced each device’s access control policies. Living in a home with 14 devices, H1B struggled with identifying and remembering which devices had access controls:

Seeing the list of all of the devices in the room, and knowing which ones he could click, and which ones he couldn’t, and which ones had to ask for permission... (H1A-Exit)

H1 suggested an interface for favorite devices (a feature supported by Vera), while H3 suggested that devices that you did not have access to would simply be hidden.

Install Barrier. We attempted to make the install process as painless as possible for our app, implementing a QR-code

passwordless public key authentication system for additional users. However, even this barrier was too much for some users — H1A did not want to go to the effort for adding their nanny (despite stating the desire to set access controls for her), and H8A did not feel confident in being able to add her mother-in-law without our guidance. As a result, these household members were either shut out of the smart home, or accessed it via other means (i.e. Amazon Echo), bypassing our prototype’s access controls and losing access to notifications.

7 Discussion

7.1 Lessons on Smart Home User Behaviors

Based on our *in situ* prototype evaluation, we surface lessons about users’ security and privacy behaviors in smart homes, including how they interact with concrete security and privacy features in practice, and how our observations of actual behavior align with user preferences identified in prior work.

Limited Usage of Access Controls. Though our participants mentioned multiple use cases for access controls in our initial interviews, such as restrictions on guests, domestic workers, and children, in practice, few of them made use of the access controls we implemented. There are several possible reasons for this. In two cases, usability was a barrier; one household was discouraged by the complexity of the access control interface, and the other by the difficulty of onboarding guests. More commonly, participants did not have a strong need to use access controls, either because they were unconcerned about restricting access to mundane devices, or that existing social norms and trust in their household checked against bad behavior. Lastly, some participants chose not to use access controls because it would interfere with other desired functionality, like occasionally allowing children remote access.

These findings suggest that while at first glance there are many user goals that could be achieved with access controls, there are only a few specific use cases that access controls are well suited for in practice, like limiting access for domestic workers. But for other use cases where users have weak or subtle preferences, access controls can be too rigid, complex, or simply not useful, even with reactive and contextual mechanisms, such as parental controls.

Importance of Social Norms. Among our study population, we observed that in circumstances where prior work has shown the potential for multi-user conflicts and privacy issues, our participants often did not experience these problems due to the norms of interpersonal behavior in their home. For example, children were trusted to follow rules, roommates respected each others’ spaces, and people were not concerned about information revealed by the smart home when it matched their household’s privacy norms. This finding suggests that in generally cooperative households, multi-user security and privacy issues may be able to be addressed in

part by cultivating good norms around usage of the smart home. We discuss this topic further below.

Acceptance of Security and Privacy Tradeoffs. As we expected from prior work [41], participants were willing to accept (multi-user) security and privacy risks posed by usage of the smart home because of the convenience and utility it provided. Participants often explicitly mentioned the tradeoff between convenience and privacy, when asked about their concerns about data privacy. H8 decided against setting up access controls (for parental controls) because the smart home would be less convenient for the household, and H1 decided against using access controls for their nanny because the setup process would be inconvenient. While this finding is not new, it re-emphasizes that when designing security and privacy features for smart homes, these features must work with, and not limit, users' primary use cases for the smart home.

7.2 Revisiting our Design Principles

In Section 4, we proposed a set of design principles which we hypothesized could help address multi-user security and privacy issues. Based on the insights provided by our evaluation and user study, we revisit these principles:

Access Control Flexibility: Important But Not a Panacea.

Our results suggest that while access controls might not be suitable for satisfying all user preferences, the flexible access control mechanisms we implemented, such as location-based access controls and per-device ownership, can help users in clear-cut use cases, like guest access. However, we also found that increasing flexibility also increases the complexity of the interface, and as we discuss below, a challenging open question remains how to support such a complex array of options in a usable and useful way.

User Agency and Respect: Dominated by Social Norms.

Contrary to our initial hypotheses, we found that our participants relied more heavily on household social norms to support user agency and minimize conflicts than the access control, notification, and device discovery features we designed in our prototype. While such norms would not exist in abusive or adversarial households, for generally cooperative households, we propose a new research and design question that we discuss further below: how can a multi-user smart home be designed to *support and leverage* positive social norms, rather than existing alongside or supplanting them?

Transparency of Smart Home Behaviors: Inconclusive.

Our results suggest that smart home transparency features did not provide significant benefits for our participants, in terms of our design principles (user agency and respect among users). Participants were generally indifferent to the information provided by the activity and discovery notifications, though some participants found them to be useful for other reasons: home security and verifying that their automations were working. However, our investigation is not sufficient to conclude that

transparency might not be valuable in other contexts, e.g., with cameras or voice assistants, or among people with more adversarial relationships. It is also possible that our implementation of transparency via notifications was not effective, and that another design, like calendar [26] or dashboard [5] interfaces, would provide different reactions.

7.3 Design Recommendations and Challenges

Based on our findings and revised design principles, we surface several design recommendations for multi-user smart home systems, particularly for platforms that can orchestrate access controls and features across all devices of the home.

Support Smart Home-Specific Access Control Needs.

Our study highlights a number of use cases for access controls that appear to be common in smart home settings, including restrictions on visitors, and different policies for different rooms. To support these use cases, we recommend that smart home platforms support the following primitives: (1) Location/proximity-based access control, for handling guests and domestic workers, as well as restricting access to media devices, (2) Time-based access control, also for guests, (3) per-device roles for private rooms, (4) and per-user roles, for limiting access to device and access control configuration.

Simplify Access Control Configuration. A system with all of the above access control mechanisms will run into serious usability challenges if it simply surfaces a large matrix of multi-dimensional per-user, per-device options. In fact, such complexity risks increasing the access gap between the smart home's primary user and others with less technical or interpersonal power. It could also put the use of access control out of reach for novice users. Moreover, complex policies could introduce errors or conflicts between access control rules.

A good first step towards simplifying smart home access control could be to use sensible defaults based on data on people's access control preferences, as suggested by He et al. [16]. However, our results suggest that individual factors, social norms, and conflicting use cases may cause household needs to diverge from these broad preferences, so it is still important to have a usable configuration interface. However, it is not clear what kind of interface would be effective in this context. In Section 7.4, we recommend that future work investigate systems for simplifying access control configuration in smart homes, such as natural language-based policy creation.

Incorporate Voice Assistants into Access Control Systems.

A major limitation of our prototype was that our access control system could be (intentionally or unintentionally) bypassed by sending a command through a voice assistant, such as the Amazon Echo. This is likewise a challenge for current smart home platforms: in platforms like SmartThings, voice assistants and other third party apps like IFTTT are given unrestricted access to smart home devices via OAuth integrations. Additionally, current voice assistants do not explicitly

perform voice recognition, so a smart home would not be able to identify who is issuing a command. In order for access controls to be consistently applied, voice assistants should support voice-based authentication, and voice assistant manufacturers should work with smart home platforms to develop a federated access control system. This is particularly important as adoption of voice assistants increases and they become a popular way to interact with smart homes.

Reduce User Onboarding Barrier. The smart home control interface still needs to be made more accessible to users. Even by our best efforts, a mobile app was too much to ask for some participants to install without our direct assistance and urging. If a smart home control system provides perfect security and privacy features that are locked up in an app that not all household members install, the benefits of these features will be limited. And in worst-case scenarios, if a household members cannot gain access to the smart home, it can enable domestic abuse by those with control. We suggest several potential approaches to address this issue:

One approach is to lower the installation barrier by making a mobile web version of control interfaces. In our experience, Web APIs were sufficient for all functionality, except for Bluetooth beacon scanning for proximity sensing — though browsers intend to implement this feature in the future [7].

Another approach is to further simplify user authentication. Our prototype required only a QR code rather than a username/password for subsequent users. We suggest exploring even more radical approaches, such as not requiring *any* traditional authentication to use the smart home, and instead granting basic smart device control functions to anyone in physical proximity (just as someone with physical access to a manual light switch can toggle it).

7.4 Directions for Future Research

Our work also suggests research questions that we encourage future work to investigate:

Study and Design for Positive Household Norms. We observed in our study that in cooperative households, social norms were effective at mitigating multi-user security and privacy issues, sometimes more so than the features we implemented in our prototype. Rather than trying to provide features that play the same role as these social norms, like location-based access controls for preventing inconsiderate use of remote access, we suggest (1) studying households that exhibit positive social norms around smart home usage and (2) designing and evaluating smart home systems that encourage the development of these norms in generally cooperative households. Based on the results of our study, we propose a few design “nudges” that could potentially instill better behaviors in smart home users.

First, rather than asking users to design access control policies around considerate usage, smart home platforms could

automatically detect commands that are potential norm violations, and then ask the user “Are you sure?”, including a reason for why the command might violate a norm. For example, this prompt could be triggered when attempting to control devices in another user’s private bedroom, or when remotely controlling devices that would impact other people physically present. Such a prompt could encourage users to think twice about disturbing others, while still allowing for seamless access if necessary.

Another type of nudge could promote user agency: during the setup of a smart home, the app could encourage the person installing the smart home to involve other occupants in the setup process, including encouraging and even guiding the setup of additional accounts and conversations about the different devices, automations, and policies that should be part of the new smart home. How to best design such a conversational guide is an interesting question for future work.

Nudges could also be designed to “scold” users for excessive trolling or other playful behavior, like rapidly flicking lights on and off. While it might be good to allow playful experimentation when the smart home is set up initially, eventually the app could rate limit these behaviors, or display a dialogue box encouraging the user to stop.

While norm-based nudges would of course not protect against users with malicious intent, our study results suggest that promoting positive norms could help reduce friction in the case of generally cooperative households, where conflicts and tension may arise from unfamiliarity with how one’s actions affect others in the smart home. Next, we discuss the challenge of designing smart homes for adversarial settings.

Investigate Designs for Adversarial Situations. Smart homes can enable or amplify harms in adversarial living situations, like in households where domestic abuse is occurring, or in homes with Airbnb-style rentals. While some of the design principles we proposed could mitigate some of these harms, such as using notifications to provide more transparency about how surveillance cameras are being used, our prototype would not provide adequate protections against other harmful actions, such as a malicious admin abusing their privileges to deny victims control of the home, or overriding protections against remote harassment that location or role-based access controls could provide. This is a very challenging problem, because some of these security and privacy features are inherently dual use: for example, admin roles and access controls may be desirable for parents to prevent children from doing harmful things, but could be used by abusers to exercise power over their victims. A critical but challenging design question for future work is how to design smart home access controls and monitoring that both protects users from abuse, but still enables benign use cases.

Study Transparency Features for Privacy-Sensitive Devices. As discussed above, a limitation of our prototype was that we could not provide activity notifications for privacy

sensitive smart home devices like voice assistants and security cameras, because of the limitations of the SmartThings API. We suspect that surfacing information about when audio and video is being recorded or viewed could change users' perceptions of the privacy risks of these devices, and could help people identify when their privacy is being violated. We propose an *in situ* evaluation of user reactions to a smart home system that notifies people if they are being recorded, or if another user views or listens to a log that they are present in.

Audio/video recording notifications could also be surfaced not just in the smart home, but at a global level with cooperation from mobile operating systems and device manufacturers. Cameras and microphones could emit Bluetooth beacon signals when they are active, so that users could receive notifications whenever they are nearby an active recording device.

Study Natural Language-based Access Control Policy Creation for Smart Homes. During our interviews, we observed that our participants were able to clearly convey their access control preferences and hypothetical policies verbally. Given that these policies are easily comprehensible in natural language, a possible way to simplify configuration is to allow users to craft policies using a natural language interface, rather than menus with drop-down lists and checkboxes. While prior work has found that direct conversion from natural language to policy is possible but imprecise [20, 31, 32], controlled natural language policy creation could be used to constrain the space of usable words and sentence structures. Using a controlled natural language approach, a possible interface could be an autocomplete-style input, which guides users through picking access control mechanisms, possible devices, users, roles, and other conditions. While this approach was found to be relatively usable in a systems administration context [17, 30], future work should evaluate whether it is usable for typical end users in a smart home setting.

Further Study of Automations and Attributions. We were not able to fully study whether notifications could help users with debugging automations, or attributing issues caused by automations and third-party apps, because of technical limitations of our prototype (specifically, that SmartThings does not surface to third-party applications the provenance of programmatic smart device actuations). Other researchers have proposed ways of preventing buggy or malicious behavior by third-party smart home integrations, such as detecting provenance [37] or contextual permission prompts for third-party apps [19]. These research contributions are technically valuable but their usability and utility have not been tested with real end users; we suggest that future work do so.

7.5 Limitations

Though an in-home user study allowed us to study how people used our prototype under realistic circumstances, this study design nevertheless comes with several limitations.

Most importantly, as discussed already, our prototype and user study focused on generally cooperative households, rather than households with adversarial relationships. Since we required consent from all participating household members, our sample is skewed towards households with sufficiently functional interpersonal relationships to agree to participate together in the study. Thus, we were unable to evaluate how our prototype would perform in an adversarial setting, nor did we gain insight into how to design for those settings.

Moreover, our protocol design involved conducting interviews with participants in a group setting, with the entire household. It is possible that participants were unwilling to reveal multi-user conflicts and privacy issues, because it would also reveal these problems to other household members.

Additionally, the devices our participants chose were generally not among the most invasive. This was due both to technical limitations (e.g., our prototype could not integrate with most security cameras using the SmartThings API), and because we gave participants the freedom to choose devices they were comfortable with. While our prototype did not interface with these more privacy sensitive devices, we still learned from participants via hypotheticals about access control grounded in their concrete experiences with our prototype and their past experiences with those devices. Future work should further consider multi-user smart home design in the face of more invasive devices.

Finally, the complexity and cost of an in-home study limited the feasible number of participating households, preventing us from drawing any quantitative conclusions from our results.

Despite these limitations, we believe our study provides valuable insights into how to design multi-user smart home security/privacy features for many (though not all) households.

8 Additional Related Work

Methodologically, our paper drew on a number of other in-home studies of smart homes, from HCI and ubiquitous computing. Most closely related to our work were the design and evaluation of a calendar-based interface for smart home control [26], and of a smart home data visualization dashboard [5]. Other in-home studies in HCI have studied how users interact with commercial smart homes in practice, like general usage patterns and usability [4, 18, 25], setup and configuration [9], and end user programming [39]. Researchers have also studied how users perceive and use privacy sensitive devices like cameras and voice assistants, both in-situ [28, 40], and in interviews or surveys with broader populations [21, 42].

In terms of the security and privacy of smart home devices and platforms, researchers have discovered vulnerabilities in the underlying protocols and technologies (e.g., [2, 15, 29, 38]) and studied the spread and behavior of the Mirai botnet that targeted IoT devices [1]. Other work has analyzed security and privacy weaknesses in smart home platforms that support third-party apps like SmartThings [11]. To address the

risks posed by apps, researchers have proposed and evaluated various defenses, including modifications to trigger-action programming platforms to limit misuse of access tokens [13], restricting apps using flow control [12], using provenance detection to identify anomalies [37], and a contextual access control system to protect against malicious third-party apps [19].

9 Conclusion

Multi-user smart homes face unique security and privacy challenges, such as supporting a wide range of access control preferences, and managing tensions and conflicts between users. Finding the design of current smart home systems to be insufficient for addressing these challenges, and recognizing the gap in knowledge around what designs can meaningfully improve end user experiences, we conducted an in-home user study to investigate possible approaches and solutions. Focusing on generally cooperative (rather than explicitly adversarial) households, we designed a smart home control interface based on design principles of access control flexibility, user agency, respect among users, and transparency of smart home behaviors. We deployed our prototype in seven households in a month-long study to evaluate our proposed design principles, and to improve our understanding of how users interact with security and privacy features in practice. Based on the findings of our user study, we provide design recommendations and identify open challenges for future research. Among our recommendations, we suggest that researchers improve the usability of smart home access controls by developing more usable configuration interfaces (such as natural language policy creation), and design smart home platforms that reduce tensions and conflicts by leveraging and scaffolding positive household norms.

Acknowledgements

We are extremely grateful to our user study participants for making this research possible, as well as our pilot study participant, Greg Akselrod. We would like to thank Christine Geeng, Ivan Evtimov, Kiron Lebeck, and Shrirang Mare for reviewing an earlier draft of this paper. We would also like to thank Tadayoshi Kohno for his feedback in the early stages of this research. We thank Sarah Mennicken for her advice on conducting in-home user studies. Lastly, we thank our anonymous reviewers and our shepherd, Sascha Fahl, for providing us valuable feedback for improving our paper. This research was supported in part by the National Science Foundation under Award CNS-1513584.

References

- [1] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou. Understanding the Mirai Botnet. In *26th USENIX Conference on Security Symposium*, 2017.
- [2] R. Baldwin. Researcher finds huge security flaws in Bluetooth locks. <https://www.engadget.com/2016/08/10/researcher-finds-huge-security-flaws-in-bluetooth-locks/>, 2016.
- [3] N. Bowles. Thermostats, Locks and Lights: Digital Tools of Domestic Abuse. <https://www.nytimes.com/2018/06/23/technology/smart-home-devices-domestic-abuse.html>, 2018.
- [4] A. B. Brush, B. Lee, R. Mahajan, S. Agarwal, S. Saroiu, and C. Dixon. Home Automation in the Wild: Challenges and Opportunities. In *SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 2115–2124, New York, NY, USA, 2011. ACM.
- [5] N. Castelli, C. Ogonowski, T. Jakobi, M. Stein, G. Stevens, and V. Wulf. What Happened in my Home?: An End-User Development Approach for Smart Home Data Visualization. In *CHI Conference on Human Factors in Computing Systems (CHI)*, 2017.
- [6] E. K. Choe, S. Consolvo, J. Jung, B. L. Harrison, S. N. Patel, and J. A. Kientz. Investigating receptiveness to sensing and inference in the home using sensor proxies. In *14th International Conference on Ubiquitous Computing (UbiComp)*, 2012.
- [7] Chromium blink-dev mailing list. Intent to Implement: Web Bluetooth Scanning. <https://groups.google.com/a/chromium.org/forum/#!topic/blink-dev/aVxGkVQ2xRk>, 2018.
- [8] N. Dell, V. Vaidyanathan, I. Medhi-Thies, E. Cutrell, and W. Thies. “Yours is better!”: Participant Response Bias in HCI. In *SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2012.
- [9] A. Demeure, S. Caffiau, E. Elias, and C. Roux. Building and Using Home Automation Systems: A Field Study. In *International Symposium on End User Development (IS-EUD)*, 2015.
- [10] Y. Elrakaiby, F. Cuppens, and N. Cuppens-Boulahia. Interactivity for Reactive Access Control. In *International Conference on Security and Cryptography (SECRYPT)*, 2008.
- [11] E. Fernandes, J. Jung, and A. Prakash. Security Analysis of Emerging Smart Home Applications. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 636–654, 2016.
- [12] E. Fernandes, J. Paupore, A. Rahmati, D. Simionato, M. Conti, and A. Prakash. FlowFence: Practical Data

- Protection for Emerging IoT Application Frameworks. In *USENIX Security Symposium*, pages 531–548, Austin, TX, 2016. USENIX Association.
- [13] E. Fernandes, A. Rahmati, J. Jung, and A. Prakash. Decentralized Action Integrity for Trigger-Action IoT Platforms. In *Network and Distributed System Security Symposium (NDSS)*, 2018.
- [14] C. Geeng and F. Roesner. Who’s In Control?: Interactions In Multi-User Smart Homes. In *CHI Conference on Human Factors in Computing Systems (CHI)*, 2019.
- [15] J. Granjal, E. Monteiro, and J. Sá Silva. Security for the Internet of Things: A Survey of Existing Protocols and Open Research Issues. *IEEE Communications Surveys Tutorials*, 17(3):1294–1312, thirdquarter 2015.
- [16] W. He, M. Golla, R. Padhi, J. Ofek, M. Dürmuth, E. Fernandes, and B. Ur. Rethinking Access Control and Authentication for the Home Internet of Things (IoT). In *USENIX Security Symposium*, 2018.
- [17] P. Inglesant, M. A. Sasse, D. W. Chadwick, and L. L. Shi. Expressions of expertness: the virtuous circle of natural language for access control policy specification. In *SOUPS*, 2008.
- [18] T. Jakobi, C. Ogonowski, N. Castelli, G. Stevens, and V. Wulf. The Catch(es) with Smart Home: Experiences of a Living Lab Field Study. In *CHI 2017*, 2017.
- [19] Y. Jia, Q. A. Chen, S. Wang, A. Rahmati, E. Fernandes, Z. M. Mao, and A. Prakash. ContextIoT: Towards Providing Contextual Integrity to Appified IoT Platforms. In *Network and Distributed System Security Symposium (NDSS)*, 2017.
- [20] H.-T. Le, D. C. Nguyen, L. C. Briand, and B. Hourte. Automated inference of access control policies for web applications. In *SACMAT*, 2015.
- [21] N. Malkin, J. Bernd, M. Johnson, and S. Egelman. “What Can’t Data Be Used For?” Privacy Expectations about Smart TVs in the US. In *European Workshop on Usable Security (Euro USEC)*, 2018.
- [22] S. Mare, L. Girvin, F. Roesner, and T. Kohno. Consumer Smart Homes: Where We Are and Where We Need to Go. In *IEEE Workshop on Mobile Computing Systems and Applications (HotMobile)*, 2019.
- [23] T. Matthews, K. O’Leary, A. Turner, M. Sleeper, J. P. Woelfer, M. Shelton, C. Manthorne, E. F. Churchill, and S. Consolvo. Stories from survivors: Privacy & security practices when coping with intimate partner abuse. In *CHI Conference on Human Factors in Computing Systems*, 2017.
- [24] M. L. Mazurek, P. F. Klemperer, R. Shay, H. Takabi, L. Bauer, and L. F. Cranor. Exploring Reactive Access Control. In *CHI ’10 Extended Abstracts on Human Factors in Computing Systems*, 2010.
- [25] S. Mennicken and E. M. Huang. Hacking the Natural Habitat: An In-the-Wild Study of Smart Homes, Their Development, and the People Who Live in Them. In *International Conference on Pervasive Computing (Pervasive)*, 2012.
- [26] S. Mennicken, D. Kim, and E. M. Huang. Integrating the Smart Home into the Digital Calendar. In *CHI Conference on Human Factors in Computing Systems (CHI)*, 2016.
- [27] C. Nandi and M. D. Ernst. Automatic Trigger Generation for Rule-based Smart Homes. In *ACM SIGPLAN Workshop on Programming Languages and Analysis for Security (PLAS)*, 2016.
- [28] A. Oulasvirta, A. Pihlajamaa, J. Perkiö, D. Ray, T. Vähäkangas, T. Hasu, N. Vainio, and P. Myllymäki. Long-term Effects of Ubiquitous Surveillance in the Home. In *14th International Conference on Ubiquitous Computing (UbiComp)*, 2012.
- [29] E. Ronen, A. Shamir, A.-O. Weingarten, and C. O’Flynn. IoT Goes Nuclear: Creating a ZigBee Chain Reaction. *IEEE Symposium on Security and Privacy*, 2017.
- [30] L. L. Shi and D. W. Chadwick. A controlled natural language interface for authoring access control policies. In *SAC*, 2011.
- [31] J. Slankas and L. A. Williams. Access control policy extraction from unconstrained natural language text. *2013 International Conference on Social Computing*, pages 435–440, 2013.
- [32] J. Slankas, X. Xiao, L. A. Williams, and T. Xie. Relation extraction for inferring access control rules from natural language artifacts. In *ACSAC*, 2014.
- [33] SmartThings Community Forums. Guest Access - Solution? <https://community.smartthings.com/t/guest-access-solution/97288/26>, 2017.
- [34] M. Surbatovich, J. Aljuraidan, L. Bauer, A. Das, and L. Jia. Some Recipes Can Do More Than Spoil Your Appetite: Analyzing the Security and Privacy Risks of IFTTT Recipes. In *26th International Conference on World Wide Web (WWW)*, 2017.
- [35] B. Ur, J. Jung, and S. Schechter. The Current State of Access Control for Smart Devices in Homes. In *Workshop on Home Usable Privacy and Security (HUPS)*. HUPS 2014, 2013.
- [36] B. Ur, J. Jung, and S. E. Schechter. Intruders Versus Intrusiveness: Teens’ and Parents’ Perspectives on Home-Entryway Surveillance. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2014.
- [37] Q. Wang, W. U. Hassan, A. M. Bates, and C. A. Gunter. Fear and Logging in the Internet of Things. In *Network and Distributed System Security Symposium (NDSS)*, 2018.
- [38] M. Wollerton. Here’s what happened when someone hacked the August Smart Lock. <https://www.cnet.com/news/august-smart-lock-hacked/>, 25 2016.

- [39] J. Woo and Y.-K. Lim. User Experience in Do-It-Yourself-Style Smart Homes. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2015.
- [40] P. Worthy, B. Matthews, and S. Viller. Trust Me: Doubts and Concerns Living with the Internet of Things. In *ACM Conference on Designing Interactive Systems (DIS)*, 2016.
- [41] E. Zeng, S. Mare, and F. Roesner. End User Security and Privacy Concerns with Smart Homes. In *Symposium on Usable Privacy and Security (SOUPS)*, 2017.
- [42] S. Zheng, N. Apthorpe, M. Chetty, and N. Feamster. User Perceptions of Smart Home IoT Privacy. *Proceedings of the ACM on Human-Computer Interaction (PACMHCI)*, 2:200:1–200:20, 2018.

Appendices

A Initial Interview Script

Control and Agency

- Who found out about the study? Who wanted to be a part of it?
- Did you set up your smart home together, or did one person take the lead?
- Do all of you have access to all smart home devices right now? If not, why?
- What are you hoping that your smart home will do for you?

Multi-user Privacy

- Have you ever unexpectedly learned anything about someone else, through the smart home?
- Can you think of ways you could “spy” on people using your smart home? Would you do it?
- Do you think it’s a good or bad thing that you can find out those things?

Transparency

- Are you having any trouble figuring out how to control your devices? Or figuring out which devices are smart?
- Has there been any confusing moments where you weren’t sure what was causing something to happen in your home? How did you figure it out?

Access Control Preferences

- Can you think of any situations where you want to restrict where people could remotely control devices from?
- Can you think of any situations where you want to restrict certain people from controlling certain devices?

General Security and Privacy Questions

- Do you have any security and privacy concerns about smart homes?
- Are there any potential security and privacy issues that you are aware of, but aren’t worried about?

B Exit Interview Script

General Usage and Control

- How did you end up using your new devices?
- Did you set up any automations?
- How involved were each of you in configuring the home? Like setting up rooms and permissions?

Notifications and Transparency

- Let’s talk about the activity notifications feature - the notifications you can get when someone turns something on or off, or trips one of your sensors. Did you have this feature on? (Why not?)
- How did you set your preferences for notifications? Why? Which devices? Proximity based or not?
- In what situations did you normally see notifications?
- Did seeing notifications provide any useful or interesting information?
- Did any notifications help you understand what your smart home was doing?
- Did you learn something about other people’s behavior that you wouldn’t have found out about without notifications?
- Did you change your behavior in your home because of the notifications?
- Were the notifications overwhelming, or not useful?
- What changes would you like made to make to this feature?
- Leaving aside the particular capabilities of our app, can you think of any situation where it would be useful to get notifications, maybe just for particular devices?

Supervisory, Reactive, and Role-based Access Control

- Let’s move onto the allowed users feature. This is the feature that lets you designate owners for each device, and have everyone else ask for permission to use it. Did you use this feature?
- If so, who was restricted? What devices and policies did you set? (block, ask, ask if not nearby)
- If not, why?
- How did you all decide on who to set restrictions on?
- In what situations did <restricted user> have to ask for permission to use a device?
- Did anyone try to circumvent restrictions on them? How?
- To blocked user: was it clear to you which devices you needed permission to use? How did you find out?
- To blocked user: How comfortable did you feel pushing the button to ask for permission?
- To blocked user: Did you change your behavior as a result of having to ask for permission?
- To admin users: How did you feel when you got notifications when someone asked for permissions?
- To blocked user: when you asked for permission, did the other person usually respond in time?

- To admin user: did you receive notifications in a timely manner? Were you able to fulfill requests?
- What changes would you like made to make to this feature?

Location-based Access Control

- Now let's talk about permissions for remote control. This is the setting where you can make people ask for permission to use a device if they aren't nearby. Were any devices restricted to remote control in a particular location? If not, why?
- How did you all decide on which devices to set restrictions on?
- In what situations did you have to ask for permission to use a device?
- Did anyone try to circumvent the restrictions on a device? How?
- When you had to ask for permission, did someone respond in time?
- When you got a permission request, did you receive a notification in a timely manner? Were you able to fulfill the request?
- Did the beacons usually accurate put you in the correct room?
- Was it clear which devices were location restricted? How did you know?
- To blocked user: How comfortable did you feel pushing the button to ask for permission?
- To blocked user: Did you change your behavior as a result of having to ask for permission?
- To admin users: How did you feel when you got notifications when someone asked for permissions?
- Did you ever use this feature to check who was home?
- What changes would you like made to make to this feature?
- Hypothetically, imagine we built an app that had every access control scheme and level of granularity you wanted - custom permission tiers, time-based access controls, proximity-based access controls, and device-level granularity. How would you set these for the different people who visit your home? (Spouse, children, guests, domestic workers?)

C Codes Used for Qualitative Analysis

- Access control - ask for permission
- Access control - complexity/discoverability
- Access control - conflicts with other goal

- Access control - desired use cases
- Access control - location-based
- Access control - not useful
- Access control - role-based
- Access control - side channel
- Access control - trust/respect each other
- Access control - unconcerned about device
- Access control - useful
- Multi-user - conflicts
- Multi-user - pranks
- Multi-user - privacy
- Multi-user - unexpected home behavior
- Notifications - checking/debugging automations
- Notifications - desired use cases
- Notifications - not noisy
- Notifications - not useful
- Notifications - privacy
- Notifications - proximity scoping
- Notifications - too noisy
- Notifications - useful
- Relationship - children
- Relationship - couples
- Relationship - domestic workers
- Relationship - guests
- Relationship - roommates
- SecPriv - Accepts risk
- SecPriv - Concern about location/proximity
- SecPriv - Concern about others
- SecPriv - Non concern
- SecPriv - Privacy concerns
- Usability - automation confusion
- Usability - complexity
- Usability - discoverability/naming
- Usability - install barrier
- Usability - need phone
- Usability - setup difficulty
- Utility - automation
- Utility - general convenience
- Utility - provides security
- Utility - remote control
- Utility - time cost