Access Right Assignment Mechanisms for Secure Home Networks

Tiffany Hyun-Jin Kim, Lujo Bauer, James Newsome, Adrian Perrig, and Jesse Walker

(Invited Paper)

Abstract: **The proliferation of advanced technologies has been altering our lifestyle and social interactions–the next frontier is the digital home. Although the future of smart homes is promising, many technical challenges must be addressed to achieve convenience and security. In this paper, we delineate the unique combination of security challenges specifically for access control and consider the challenges of how to simply and securely assign access control policies to visitors for home devices and resources. We present a set of intuitive access control policies and suggest four access control settings based on our in-person interview results. Furthermore, we propose the automated Clairvoyant access right assignment (CARA) mechanism that utilizes home owners' social relationship to automatically deduce to which class a visitor belongs. The combination of CARA and the suggested mapping provides a promising first step for home policy assignment such that non-expert home owners can let visitors use their home network with confidence. We anticipate that future research can build on our proposed mechanisms to provide confidence to non-expert home owners for letting visitors use their home network.**

Index Terms: **Access control, future home networks, privacy, security, usability.**

I. INTRODUCTION

We are in year 2020. Alice and Bob are working parents who recently purchased a "smart" McMansion house which is equipped with the most recent appliances, devices and communication technology for networking all home devices. This house provides a data storage system storing entertainment data (e.g., games, movies, music), family photos, personal files (diary), highly sensitive information (tax records), etc. In this environment, the house is the computer, most surfaces can act as displays (including walls and appliance surfaces), and resources can be accessed from anywhere. Also, all individuals have smart phones and they can interact with resources in the house through wireless communication. We anticipate the following scenarios that will motivate the research challenges we will encounter and address in our research.

Alice and Bob have three children, Carol (15), David (12), and Evan (5). The children have a need to access most of the functionality available through the home network, from using printers and file servers to controlling the heating and cooling and giving visiting friends access to the network. Despite their general high-level (super-user) access, the children are restricted from a small subset of activities, such as accessing the family's tax returns on the home file server or allowing complete strangers to gain Internet access through the home network.

When Carol introduced her boyfriend Frank to her parents, Alice and Bob were put off by his disheveled hair style and threatening skull tattoos and consequently gave him only very limited access to the home network. Over time, however, they came to increasingly trust him, and needed to adjust the security policies they initially put in place from a default very low level of trust to a gradually higher level.

David has a friend, George, who visits often and the boys play computer games. Since George visits frequently, he has access to many home network resources, which he needs to play the games. George, however, has a crush on Carol, and is eager to download family photos and videos with her as well as spy on Carol's personal diary.

Helen is the nanny who takes care of Evan during the days when the other family members are working or at school. Helen has a curious nature; she loves browsing through other people's information. Helen also has had a sequence of questionable boyfriends. Hence, Alice is concerned that Helen trusts people too easily and would delegate access rights to strangers.

When Alice's parents visit, they stay at the house for weeks at a time. Bob is always intimidated by the inquisitive nature of Alice's mother Irene; he is concerned about Irene's eagerness to view the tax returns and to inspect the fridge log to monitor Bob's beer consumption. Bob would like to restrict Irene's access rights, but is also concerned that Irene may interpret his actions as distrusting her.

Technology trends. The future smart home vision that we envision is enabled by a number of technology trends:

- **User interfaces (UIs) for "everything."** As in Mark Weiser's vision, "invisible" computers and interfaces (i.e., ease of use is so effective that one does not notice the computer) will transcend most objects we interact with [1], and appliances will have built-in computers, UIs (display, keyboard), and/or RFID tags.
- **Network communication.** Objects with computing capabilities will also connect to the home network and the Internet.

Manuscript received October 27, 2010; approved for publication by Heejo Lee, JCN Editor, March 03, 2011.

This research was supported in part by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, grant CNS-0627357 from the National Science Foundation, and by gifts from Bosch and Intel. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of ARO, Bosch, CMU, Intel, NSF, or the U.S. Government or any of its agencies.

T. H.-J. Kim, L. Bauer, J. Newsome, and A. Perrig are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA, email: {hyunjin1, lbauer, jnewsome, adrian}@ece.cmu.edu.

J. Walker is with the Intel Corporation, Hillsboro, OR 97214, USA, email: jesse.walker@intel.com.

Network communication will enable remote device operation and management.

- **Digital media.** Media will continue transitioning from physical to purely digital. Examples include MP3 files, Netflix movies, Kindle eBooks, and photos on Flickr.
- **Smartphones.** Smartphones will become universal UIs to control devices in a smart home. Compared to year 2009, smartphone sales grew 96% and smartphones accounted for 19.3% of overall mobile phone sales in Q3 of 2010.[1] In the foreseeable future the majority of phones will be smart phones, and users are already developing home control applications on smart phones.
- **Smart meters & grids.** Smart meters and grids reduce costs by enabling power companies to use demand-response mechanisms. This makes it possible to manage electricity consumption in response to supply conditions (e.g., market prices).
- **Wireless medical devices.** Many health-care devices are becoming portable and wireless to enable real-time monitoring by doctors.

These trends will fundamentally alter our living style and the way we interact with our home. Indeed, there already is a cross-industry organization of leading consumer electronics, computing and mobile device companies called digital living network alliance that enables digital content (e.g., photos, music) to be shared among devices that belong to the same network (e.g., laptops, mobile phones).[2] For many new technologies, new features drive adoption, and unfortunately, security and privacy issues are often left to be addressed later. However, a challenge is to build smart homes that are both convenient and secure. In this paper, we consider how to address the security issues of access control management in such an environment when sharing resources while minimizing user involvement.

Security issues and challenges. Consider, for example, that the home will have a plethora of microphones and cameras that can be remotely activated; sensitive data such as health and financial information will be accessible from anywhere; records of viewing and reading habits, personal photos, videos, and diaries will all be available digitally; implanted medical devices can be remotely controlled by health care providers and interact with medical databases. In this context, computer security breaches will not only compromise individuals' and families' privacy to an even greater degree than ever before, but can also easily cause direct physical harm, all in the "comfort" of one's own home.

The fundamental challenge that we focus on is how to control access in this environment—essentially, how to enable home users to manage access-control policies for everyone who visits their homes, including family members, friends, visitors (e.g., repairman, housekeeper, accountant), as well as emergency-related personnel (e.g., first responder, doctor). The central issues in this space revolve around the complexity and diversity of the resources, the diversity of the subjects, the low sophistication of the administrators, and the social context.

It is clear that existing access policy control mechanisms are woefully inadequate in this environment: discretionary access control mechanisms do not scale to the complexity of homes (it would be impractical to set access rights to hundreds of resources for each visitor), and corporate access control management systems require professional administrators.

Our approach. In this setting, our goal then is to integrate smart appliances into home networks that can guarantee security and privacy to non-expert home owners. Our approaches are: 1) By translating current home usage metaphors, establish home access control policies that are intuitive to use by mimicking the way current home devices are used by visitors and family members, 2) enable users to establish basic access policy profiles, and 3) to further simplify the task of access right assignment, use an owner's social network to automatically assign users to a basic access policy profile.

Contributions. This paper makes the following contributions:

- Based on our initial work [2], we enumerate the series of challenges that makes the access control management of the digital home a unique and particularly difficult task. Although some of the individual challenges may appear in other contexts, the home environment presents a unique combination of challenges.
- We propose a framework for reasoning about and implementing access control in future home networks. Our framework has the following key features:
  - Access control permissions and policies are organized in three orthogonal categories: Presence, logging, and asking for permission. Home owners may choose one policy from any of the combinations of these three dimensions, and assign it to a visitor.
  - Access control policies are grouped in four settings: Full control, restricted control, partial control, and minimal control.
  - We propose a baseline mapping between the set of access control policies and the groups of the access control settings. Our approach is motivated by a preliminary user study, and by adapting our suggestions, home owners can simplify the policy assignment procedure.
- We propose an automated access right assignment for visitors based on home owners' social networking data. Our automated access policy assignment algorithm provides visitors with closer social relationship to the home owner with higher access rights on home resources.
- We evaluate the completeness of the set of access control policies and the access control policy settings by conducting a user study. Also, using home owners' cellular phone usage information as the social networking data, we verify the accuracy of our automated policy assignment mechanism.

II. BACKGROUND: OUR HOME NETWORK MODEL

Since our work is based on home networks, we describe the controller device that coordinates and controls wireless communications among smart devices in the network in this section. This controller device has specific features and responsibilities; it is universally trusted and is capable of creating keys for secure communication. Moreover, the controller can set the access control rules of any device in the network.

A home owner can initiate the home network simply by purchasing a main controller device, such as Nokia home control

[1] [Online]. Available: http://www.gartner.com/it/page.jsp?id=1466313

[2] [Online]. Available: http://www.dlna.org

center[3] which will function as the root of the network. This controller device may carry standard factory-defined access control settings by default as a baseline. An example of the default assignment categories may be "full control," "restricted control," "partial control," and "minimal control."

The owner pairs her mobile device with the main controller and assigns the "full control" setting to herself. In this way, the main controller recognizes the owner's mobile device with complete access privileges. Pairing a mobile device with the controller may be challenging for a non-expert user. However, such a pairing procedure can be configured by the owner herself with detailed instructions. Secure pairing can be accomplished through location-limited channels such as infrared or near-field communication (NFC), visual channels, or a wired connection [3], [4]. Secure pairing is used to prevent man-in-the-middle attacks by a malicious neighbor, for example. An alternative solution is for an external technician to set up the entire network. An owner can use similar methods to add or remove devices to the home network, though this is largely orthogonal to the problem we study in this paper.

We assume that an access control infrastructure, e.g., a trust management system like KeyNote [5], [6] or PolicyMaker [7], is in place to implement the access control policies. The assumption is that the owner's and visitor's mobile devices (i.e., their smartphones) can set up a secure connection and that the owner can assign access rights to the visitor's device. We further assume that all mobile devices are also uniquely identifiable, for example, through a public-key certificate or secure hardware, such as a trusted platform module (TPM) chip [8].

A. User Lifecycle

In a typical home, new people enter and current ones leave on a constant basis. Enrollment and revocation of people who live at the house are not a major challenge, because these events are infrequent. A greater challenge are visitors who arrive and leave frequently. We classify the interactions between the home owner and visitors into the following four categories.

User enrollment. The owner can give guests and friends access to devices in her home network by providing appropriate access right assignments. A step that needs to take place before access rights assignment, however, is the establishment of a name-to-key binding for the guests and friends. This binding links a person's real-life identity to a digital identity to which access rights can be assigned.

Access rights assignment. Based on the propinquity of the visitor to the owner, the owner may provide different access control policies. For example, the owner may allow her closest friend to use majority of resources offered by the owner's home network except her tax files. On the other hand, the owner may assign an access control policy to a one-time visitor, such as an electrician, who will then be able to use resources only to complete the task at hand. This procedure can be done manually whenever guests and friends visit. However, users may find such access right assignment cumbersome, and automated configuration of access right policies is a challenge in these environments. For example, if the visitor is the family accountant, she can access the tax records but not the family photos; but if the grandparents are visiting, they can access the family photos but not the tax records. In this paper, we study how to enable non-expert users to reliably and safely assign access policies.

Access rights revocation. The owner may need to revoke access rights for some visitors, possibly because they have violated rules (e.g., by secretly reading the owner's personal diary), or because they are no longer associated with the owner (e.g., an electrician who has finished fixing a broken switch).

User revocation. In addition to revoking access rights that are no longer appropriate, it may be necessary to remove the name-to-key bindings that link users' real-life identities to their digital counterparts. This may happen when the home owner is certain that a particular visitor will never return to the owner's home or when there is reason to suspect that a digital identity has been compromised (e.g., after detecting leakage of a user's private key).

[3][Online]. Available: http://smarthomepartnering.com

III. PROBLEM DEFINITION AND THREAT MODEL

Establishing a home network is relatively straightforward, but a core challenge is how to enable non-expert users to safely set home access control policies. In this section, we present a concise problem definition and a threat model.

A. Problem Definition

Our central goal is to protect the resources in a home network environment against unauthorized use. More specifically, we intend to protect against misuse by visitors, as we assume that current security mechanisms can protect against malicious outsiders (e.g., we do not address key management). In particular, we aim to provide a mechanism to assist home owners in giving their visitors access to particular devices or resources within their homes. A key goal is for the mechanism to be as easy to use as possible so as to be accessible to non-experts and to generally place minimal burden on users.

An access control management mechanism should provide the following security properties:

- Secrecy and privacy of personal information (protect against undesired disclosure of data),
- integrity of personal information (protect against undesired alteration or loss of data),
- availability of resources (prevent denial-of-service (DoS) attacks against resources),
- allow only permitted accesses (prevent against misuse of devices to cause annoyance, disturbance, physical damage, or economic harm).

B. Threat Model

Our adversary model is a visitor who receives unintended access privileges from some principal in the system and misuses them. More specifically, we try to guard against a visitor who receives more permissive access rights than what the home owner wishes to grant. For example, an honest but curious visitor could attempt to read sensitive information, perform unwanted alterations to existing data, or overuse devices beyond reasonable limit (i.e., printing an entire photo album on the owner's color printer). Also, the visitor could perform disturbing operations

on the home network after he leaves, for example by playing loud music at night or shutting off the home security system.

Although other attacks such as external attacks on the communication channel [9] or device compromise are important, we focus in this work solely on access control.

IV. UNIQUE COMBINATION OF CHALLENGES

Despite the plethora of research in access control, we believe that no existing solution adequately addresses the unique set of challenges posed by home environments. Discretionary access-control mechanisms do not usably scale to the complexity of homes; it would be impractical to set access rights to hundreds of resources for each visitor. Access-control systems used in corporate environments require professional administrators. While some researchers have created tools to help users create access-control policies (i.e., SPARCLE policy workbench [10], expandable grid [11]), these tools target more constrained environments and more skilled (though still non-expert) users than will characterize the future digital home. In this section, we elaborate on specific challenges that secure home access assignment systems encounter.

No dedicated expert administrator. The typical home user lacks both the patience and the expertise required of an administrator in a corporate access control system. For example, even technologically savvy Firefox 2 users ignore an expired certificate warning from their banking websites [12]. A typical home user is unlikely to spend much time learning complex interfaces or performing tasks such as assigning access rights, auditing current policies, or auditing the access logs.

Mixed ownership. In many homes, no single person owns all devices, but each household member owns a subset of devices. Also, many shared devices exist without a single clear owner. Consequently, some devices may lack an access policy, while others have inconsistent policies.

Complexity of home environments. The number and diversity of devices and resources in homes causes tremendous complexity for access control mechanisms. For example, homes have typical appliances (washer, fridge), storage devices (for music, videos, photos, files), network-related devices (wireless router, femto cell), safety devices (smoke/gas detectors, alarms), etc. Home environments are further complicated by the high dimensional types of resources that each device supports. For instance, a portable music player is no longer used just to store and listen to music—it is also used as a storage device (contact information, videos, photos, documents) and as a scheduler. Furthermore, data adds one more layer of complexity. On a storage device (i.e., desktop computer) that is shared by house members, for example, users may store sensitive personal data along with non-sensitive data that they may want to share with others.

Diversity of visiting parties. The types people who visit homes and need access to home resources is diverse, ranging from family members and relatives, friends and neighbors to service workers, utility company, first responders (law enforcement, fire fighters), health care providers, and elderly care providers. Each party requires different access to home resources, yet generating a specific access control policy for each party under all circumstances is cumbersome.

Multiple uncoordinated administrators. In homes with multiple members, a single master administrator for the home network is not sufficient for maintenance. In case the one and only administrator is away from home, there must be an alternative administrator who knows how to manage and update the access control policies; for example, an electrician needs to access the master light control system when the master administrator, who can only change the access policies for the light control system, is on business travel. Hence, it is necessary that more than one (if not all) members of the household should be able to manage access control mechanisms.

On the other hand, only trusted people should be able to change the access control configuration. For example, small children should not be able to control the access control functions for the main security system such that they cannot grant burglars (who may approach children in a friendly manner) access to home devices.

Differences in administrator preferences. Some owners want a high level of security and privacy and do not mind high management overhead while others may be trusting and prefer low administration overhead. The level of convenience desired or disturbance tolerated can also vary. Balancing the security, privacy, and the level of convenience for different users is a significant challenge.

Social context: Distrust revelation problem. Users may not want to admit that a visitor is untrusted. As a result, the usually invisible aspect of trustworthiness becomes visible through the home access control policy. A visitor who considers himself as a close friend to the home owner may become upset to learn that he is only granted the minimum access level. Such situations may put social pressure on the home owner to provide looser access controls to avoid revealing his distrust.

V. ACCESS CONTROL POLICIES IN THE HOME

A significant aspect of the problem of securing the digital home is providing users with convenient yet trustworthy mechanisms for specifying and managing access control policy. Studies have suggested that users have varied and complex access control needs (e.g., [13]). At the same time, experience teaches us that complex policies typically cannot be adequately managed by end users, especially by non-expert users. We conducted a small user study to preliminarily determine the specific access control needs of users with respect to the future digital home (subsection V-A). We found that home users wish to restrict access to resources within their home via a small set of high-level constraints (subsection V-B). Based on the results of the study, we propose that creating several sets of policies and assigning users to these sets may meet the needs of most home users (subsection V-C).

A. User Study

We conducted a small-scale interview study to observe users' concerns related to access control and desired policies. We recruited 20 people (8 males and 12 females) within the age range of 20 to 60 years old through Craigslist and personal contacts. We asked each participant to list 8 people with whom they interact on an at least semi-regular basis. We also asked each partici-

pant to consider electronics and appliances in their future home. We then sought information about the access policies that they would set on those devices to restrict their use by the 8 contacts. More specifically, we asked various questions related to how much participants would allow each contact to access home appliances and how much they would be concerned if they violate specified access rights. To prepare participants, we mentioned various instances of the policies we describe in subsection V-B, and asked them to suggest new policies when our initial ones didn't meet their needs. For example, we asked questions about how the participant would assign access policies for the main entrance, such as "would you allow person X to unlock your door and enter the house?", "would you feel comfortable to let person X unlock the door when you are not present?", or "if the door lock keeps a record of who has operated it and you can check the record, would you allow person X to unlock the door?"

While conducting this user study, we were able to validate some of the challenges as mentioned in Section IV. We observed that the participants (mostly the heads of their households) were not technical experts. Also, the participants listed diverse devices when we asked for a list of all devices for their future home, and provided various types of people as potential visitors. The participants responded that they would be concerned if the access policy assignments were revealed to the visitors.

Among the observations we make based on the data gathered in our study are the following two. First, the three types of policies that we presented users with (subsection V-B) were sufficient to capture users' desired policies. Users made use of all three, and did not propose any others when given the opportunity to do so. Second, we observe that users tend to create fixed sets of access control policies, and assign a particular set to visitors based on the duration of their relationship and the level of trust (subsection V-C).

B. Policy Constraints

To mimic access control policies in current homes, the future digital home will need to support richer policies than simply allowing or denying access to specific resources. We propose three orthogonal dimensions for naturally constraining access control policies: Presence, logging, and asking for permission.

Presence. Many current home devices require physical presence to operate, i.e., a user must be *inside* the house to gain access. Light switches fall into this category. Although in future homes wireless control of resources will be pervasive, we would like to preserve this property of requiring physical presence. This can be accomplished with two kinds of constraints: *User presence* and *owner and user presence*.

For policies constrained by *user presence*, denoted as $\mathtt{P_U}$, the home owner allows the visitor to use the home electronics and appliances under one condition: The visitor must be physically present near the device. This policy may be the simplest that non-expert home owners may use for their home devices since any visitor may use devices as needed without bothering the owners; however, this type of policy is the most vulnerable in terms of secrecy and integrity properties, since a malicious visitor could potentially access secret information or alter information while they are near a storage device. This policy is ideal for physical devices such as a light switch, which can be operated while the visitor is in the room, and are not vulnerable to secrecy or integrity violations.

For the *owner and user present* access control policy, denoted as $\mathtt{P_{OU}}$, we additionally require that the owner of the resource is physically present. For some resources, it is obvious when the resource is accessed because of noticeable artifacts of operation, e.g., the sound made by a printer. For these devices, a natural policy is to enable the access when both the owner and user are physically present. This policy is commonly used today, as visitors can usually freely use visible resources when the owner is in the same room, under the assumption that the owner would warn them if they attempt to perform an unauthorized action, either accessing unauthorized resources or overusing them beyond a reasonable limit.

Logging. We envision that future home devices will record accesses. A *permitted with logging* policy, denoted as $\mathtt{P_L}$, requires devices to maintain detailed audit logs. Rarely accessed devices may even proactively notify their owners of accesses, e.g., via a text message. This policy assumes that users are generally aware that accesses of all devices are logged. Such logging could deter visitors from making unauthorized accesses since they are likely to be discovered by the owner. The current equivalent of this policy is a security camera that watches a resource. The log entries may be prioritized based on the importance of events such that users can easily review the logs when necessary. Correctly prioritizing the entries with illegitimate accesses while preventing the entries with legitimate accesses is yet another challenge.

With logging-based policies, a user may pretend that a malicious access was inadvertent. For example, a visitor may blame an access of a tax file on a home storage server on an overly aggressive virus scanner on the visitor's mobile device. Consequently, logging-based access control should be used for resources where such inadvertent access is implausible.

Asking for permission. Sometimes it is unclear how much access to provide to visitors. Instead of enumerating exactly all access rights, we propose that lazy evaluation is appropriate in some circumstances—the owner is contacted whenever visitors attempt to use a particular resource. We call this policy *ask for permission* and denote it with $\mathtt{P_A}$. In this manner, the owner knows exactly who is trying to use which device in her home. On the other hand, the owner may be overwhelmed with queries when several guests attempt to use resources. The current equivalent for this policy is that polite visitors would ask the owner if they are allowed to open a fancy box on a shelf, for example. The length for which access is granted may vary: the owner may grant one-time access or permit access for a specific interval. Similarly, the number of allowed uses may vary to prevent visitors from overusing any devices/resources.

Hybrid policies. The three orthogonal policy constraints can be combined. Fig. 1 demonstrates the space of policies spanned by these three orthogonal policies. For example, a policy $\mathtt{P_{UA}}$ will require user presence and ask for permission.

We denote the *always deny* policy with $\mathtt{P_X}$. For some devices or resources, owners may want to deny any access by visitors. Devices containing private information, such as tax records or a personal diary, are examples. This policy must be defined explicitly by the owner, but guarantees that the resource will be protected from unauthorized accesses by others.

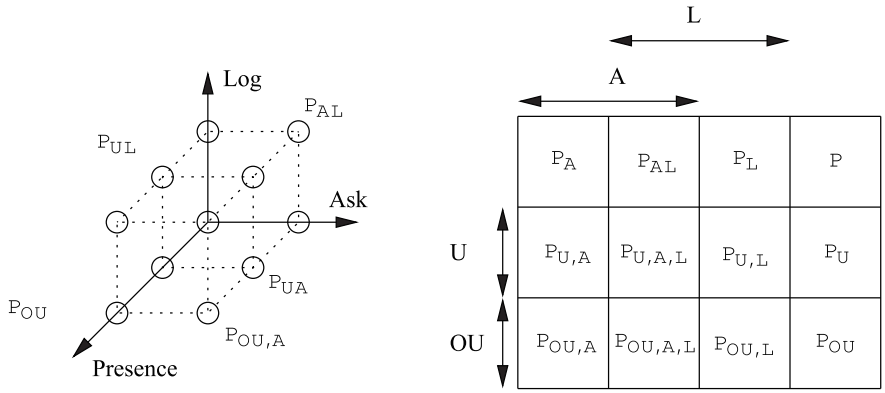

Fig. 1. Three orthogonal dimensions of our access policies, where the "presence" dimension's first point indicates "user present (U)" and the second point indicates "owner and user present (OU)." The policy corresponding to the origin (P) indicates full access without any restrictions. We illustrate five points in this space, for example, P_{UL} means that the user must be present and logging for all accesses is turned on.

C. Groups of Policies

A home owner may have a unique personal relationship with each visitor, and would hence wish to assign to that visitor a distinct set of access policies. Unfortunately, this would likely require a lot of effort.

Although studies find that categorizing all visitors into a small set of groups is unlikely, such a classification with respect to access control settings may capture *most* visitors [14]. From our user study, we observe that participants use a fixed set of categories of access control policies and assign each visitor to one of them. Such assignment is based on the length and closeness of the relationship. For example, home owners do not mind if people such as close family members and relatives open the main entrance from outside when the owners are not present; however, they would mind if people with whom they spend less time and trust less (e.g., neighbors) did so.

Based on the fine-grained responses, we were able to group access control policies into four common settings.

- **Full control**: A user is given complete control over and full access to all devices and resources. It may be assigned to owners, close relatives, and members of the household.
- **Restricted control**: Users assigned to this group of policies have full access to all devices besides the entertainment system and the security system. This group of policies may be assigned to teenagers in the household.
- **Partial control**: A user assigned to this group receives full access permissions over selected public devices that can be easily shared with others, such as a TV. This policy may be for people other than household members with whom the owner feels comfortable and whom the owner trusts.
- **Minimal control**: This setting is the most restrictive, and is granted to acquaintances or visitors who are not close friends.

Based on our study, we derive a set of specific policies with which each of these four groups could be instantiated; we show this result in Table 1.

D. Sample Real-World Integration

As mentioned in Section II, devices and resources can be manufactured with standard access right assignments by default. We suggest that devices and resources are manufactured with both basic access policies and the groups of policies as described in subsections V-B and V-C. We further suggest that devices are outfitted by the manufacturer to be able to support our suggested policy assignments as shown in Table 1. Such preloading of suggested policy assignments during manufacturing time can simplify home owners' task; instead of assigning specific policy for each and every device per visitor, they now only need to decide which of the four control settings the visitor belongs to. Then, the mapping from the control setting to basic policies for all devices and resources is automatically configured with pre-loaded suggested policy assignments.

It is possible that home owners are not satisfied with preloaded set of basic access policies, the control settings, and the suggested access policy assignments. Consequently, we suggest that devices and resources allow home owners to change policies manually; home owners can not only create new policies, new control settings, and new policy assignments, they can also modify the pre-loaded assignments that we suggest.

Table 1. Suggested basic access policy assignments for potential home devices and access control settings.

Device/resource group	Full	Restricted	Partial	Minimal
Personal laptop computer	P_U	P_U	P_A	P_A P_X
Personal file (tax/diary)				
Internet			P_U	P_A P_{OU}
Home storage (photos, music)				
Personal file storage (USB)			P_A P_L	P_A P_X
Surveillance camera				
Motion detector			P_A	
Home telephone (call log)			P_U	P_A
TV/DVR/game		P_L		P_{OU}
Digital photo frame		P_U		
Printer				
Washer/dryer/dishwasher				
Smart fridge (camera inside)				P_A P_{OU}
HVAC			P_{OU}	
Door lock			P_L	P_X P_A
Window lock				
Home security controller		P_{OU}	P_X	P_X
Room sensors (temp./light/humidity)		P_U		
Light switch			P_U	P_U
Window blind actuator				

VI. CARA: CLAIRVOYANT ACCESS RIGHT ASSIGNMENT

In this section, we present the Clairvoyant access right assignment (CARA) protocol to simplify assignment of access control rights for new people who enter the home computing environment. As we describe in Section V, an owner needs to classify a new user into a control setting, and the user inherits these access rights (as shown in Table 1). Although that operation is quick, we still aim to automate this manual assignment for the following reasons:

- *Social aspects* would force owners to provide users with too many permissions. For example, consider a casual friend that the owner puts into the "minimal control" setting instead of the "partial control" category. The visitor may contend "am I not a close friend of yours?" Automated policy assignment would circumvent this issue, as the system would automatically assign the distant friend to visitor without the owner's intervention.
- A *large number of simultaneous new users* would also put a burden on the owner, to decide for each user which category to assign to. Automated assignment would also considerably simplify the owner's tasks in this setting as well.

For this purpose, we propose CARA, a system to automatically suggest access policy assignments based on the social relationship. We observe that social networks are good indicators of the social relationship between individuals, and we thus propose to leverage social network information to determine friendship categories. As social networking information may not always be easily available, we further propose to use phone log information as we describe in this section.

A. Automated Policy Inference Using Social Networks

We observe that people trust close friends more than they trust casual acquaintances. This implies that people feel more comfortable in granting higher access rights to close friends and only minimal access rights to others. In other words, it is possible to correlate the access right level with the trust level, which can be derived from the social relationship.

We explore the social relationship among people to derive access control policies. In particular, CARA utilizes the home owner's social network such that individuals having a social relationship with a home owner can be classified into appropriate control settings based on the propinquity to the owners; CARA would categorize guests who are closer to the owner into the control setting that can access devices and resources with more privileged access rights. There are various tools from which the owner's social networking information can be retrieved. Some prominent examples are web-based contents (e.g., Facebook, blog), and instant messaging clients (e.g., Windows live messenger, AOL instant messenger (AIM), Google Talk).

Online social networks and instant messaging clients can be used to infer social relationships; the owner's social network can be retrieved based on the amount of interactions (i.e., history of overall interactions, duration of each interaction) and the rate of interactions (i.e., positive or negative interactions) with people in the websites and messengers. Consequently, people who have been interacting positively with the owner for longer amount of time would have a higher level of trust than people who have been interacting for shorter amount of time.

After CARA infers home owner's social relationship with a visitor, CARA automatically suggests the appropriate control setting for that visitor. A possibility exists that the automation process causes false alarms when assigning a control setting to a visitor. More specifically, mistakenly assigning a lower priviledge to a visitor can be mitigated without causing further damage since the visitor would request the owner to elevate her settings. On the other hand, in case the automation misclassifies a visitor and provides more permissive access rights than what the owner wants to grant, the visitor may misuse the privileges to access sensitive information, for example. In order to mitigate such an undesirable case while handling majority of assignments, we suggest that CARA assigns minimal privileges to visitors. For example, for those visitors that are analyzed to be closely related to the owner, CARA suggests that these visitors may be categorized into the "partial control" setting, and for those people that are analyzed to be distantly related to the owner, CARA suggests that they may be categorized into the "minimal control" setting. The home owner verifies and confirms CARA's suggestion if she agrees. In case she does not agree, she can manually assign the visitor to the appropriate control setting.

B. Social Relationship Approximation Using Phone Records

The CARA model as described in subsection VI-A can make a precise suggestion of the control setting into which a visitor is classified under the assumption that complete and accurate trust information can be deduced solely based on the social relationship. Unfortunately, it is quite challenging to precisely deduce the trust relationship from the owner's social networking information due to variability in communication patterns. For example, the owner who chats a lot with a co-worker for business discussion does not necessarily trust the co-worker. Moreover, Mom, who is a computer-illiterate and hence never chats online, may be the most trusted person. As a result, it is possible that CARA makes a suggestion for a visitor which may lead to provide too much access than what the owner wishes to provide.

In this section, we propose a CARA system that approximates the social relationship based on the home owner's phone records. The mobile handset of the owner maintains a list of contacts that the owner has been conversing with, as well as the logs of conversations; the handset logs when the owner makes/receives calls and how long each call lasts. Such information can be used to derive the owner's social relationships, because a person tends to speak longer and at particular times (e.g., at night time) with close friends than with someone who is not closely related. This implies that the assignment policies can be automated to suggest higher access rights to those close friends that the owner has been contacting frequently over time. Moreover, people that the owner contacts the most are always logged on top of the call log list; hence, the call log of the owner's handset is suitable to retrieve her social relationships.

SMS message logs are another good indicator for retrieving the owner's social network because a person often sends more text-messages to closer friends. Moreover, context analysis of the logged messages, which are either sent to or received from the owner's contacts, can further enhance the accuracy of the

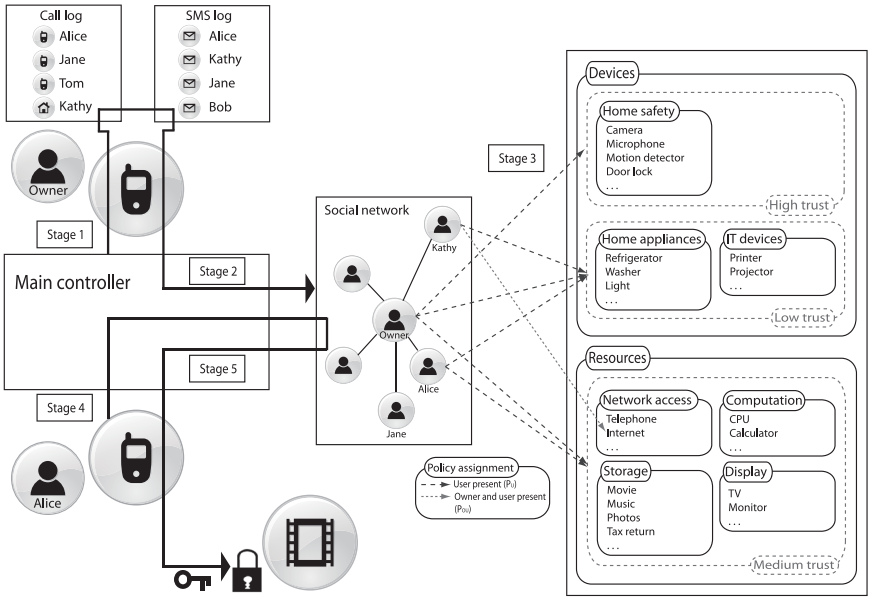

Fig. 2. A flow diagram of the automated control policy assignment procedure.

owner's social relationship analysis. For example, people tend to use special character sequences, such as emoticons, abbreviated words, and endearments, to closer friends. As a result, using both phone logs and message logs may enhance the accuracy of retrieving the handset owner's social network.

In order to reduce the misclassification error, CARA makes a suggestion for a particular visitor in terms of his trustworthiness; the visitor could be classified as either trusted or unreliable, and depending on the owner's preference, the visitor would receive the default control setting within the trusted or unreliable classifications.

C. Example: Automated Control Policy Configuration

In this section, we delineate the complete process of CARA's automated suggestion of access policy assignments using an example depicted in Fig. 2.

The home owner sets up a home network with a main controller which runs CARA to suggest access control policies for visitors. When the owner first sets up the network, CARA inspects her call and message logs from her mobile phone (stage 1 in Fig. 2), and constructs her social network as mentioned above (stage 2). Then, according to the closeness of individuals in the owner's social network, CARA suggests appropriate access control policies as described in the previous section (stage 3). The analysis procedure repeats from time to time such that CARA deduces more precise social network of the owner.

One day, the owner's friend, Alice, visits the owner to watch a movie together. Unfortunately, the owner is in the middle of an important phone call which will take for another half an hour. While waiting for the owner, Alice decides to upload the movie from her mobile device. In order to access the owner's movie storage, Alice's mobile phone sends a request to the main controller for the permission to access the movie storage (stage 4). When CARA recognizes that Alice is one of the owner's close friends who can potentially be classified as trusted and gain the permission to access the movie storage with the *logging* policy, CARA asks Alice for confirmation. Once Alice approves CARA's suggestion, CARA responds to Alice's mobile phone with an approval such that she can now upload the movie to the owner's storage (stage 5). Thanks to CARA Alice does not need to interrupt the owner to request for the storage access permission and the owner was able to focus on the important phone call without getting her guest bored.

VII. EVALUATION WITH PHONE RECORDS

Based on the previous month's cellular phone logs that 20 participants in the interview study (as mentioned in subsection V-A) agreed to provide, we were able to analyzed whom they have contacted. More specifically, we analyzed the duration and frequency of the calls to each contact. We also analyzed the time of the day (e.g., night time vs. day time) and the day of the week (e.g., weekends vs. weekdays) when calls were made to each contact. Based on this information, we formulated an algorithm that would deduce the participants' social relationship.

Our decision algorithm is based on the combination of regularity scores as described in Table 2. Each contact in the participant's phone bill (or phone log) receives certain scores if particular criteria are satisfied, and based on the threshold bounds for each class of users, the aggregated score allots the contact into either trusted or untrusted classification.

In order to assess the validity of our automated access policy suggestion mechanism using the phone usage to deduce the social relationship, we have conducted an online survey. We recruited 60 participants through Amazon mechanical turk to run the online version of our study. We were able to verify that phone records were a good indication to deduce social relationships in terms of trustworthiness; we were able to distinguish trusted and untrusted people with reasonable accuracy. As Fig. 3

Table 2. Social relationship decision algorithm based on the cellular phone usage.

Criteria (1 month)	Condition per contact	Score
Percentage of total calls	10%	2
	$\geq 1\%$ but $< 10\%$	1
	$< 1\%$	0
Duration of the longest call	≥ 15 minutes	1
	< 15 minutes	0
Calling frequency	At least twice a week	1
	At most once a week	0
Calling time (for any call)	Weekdays at 9pm to 1am or weekends at anytime	1
	Otherwise	0

shows, for a threshold value of 3, we encounter that 1 out of 19 untrusted people would have been classified as trusted (because that particular user frequently calls her untrusted boss), and 29% of trusted people have a score less than 3 because they don't call their close friends sufficiently often.

VIII. SECURITY ANALYSIS

We have shown that CARA performs accurate classification in most cases. We next consider some methods by which an attacker may attempt to gain undue privileges by abusing CARA's classification system, and some simple extensions that can be used to help prevent and mitigate such attacks.

A. *False Friend Attacks*

A user may attempt to appear to CARA to have a closer relationship to the owner by increasing his communication with the owner, thereby manipulating CARA's decision data. The attacker can take advantage of this misclassification to gain undue access to the owner's home network.

Mitigation 1: Asymmetric scoring. CARA's classification of a given user should be based on how much the owner trusts that user, and not vice versa. Such relationships are not always symmetrical. This can be captured in CARA by distinguishing between communication initiated by the owner versus communication initiated by the user. For example, a user who initiates many communications to the owner, but to whom the owner never or rarely replies, is likely to be relatively untrusted.

Hence, the simplest way to address the false friend attack is to only count communication that originates from the owner. This strategy is slightly tricky to implement in the current version of CARA, since we examine telephone communication, which is bidirectional. For example, the owner's mother may call him frequently, but rarely vice versa. We do partially address this issue by giving longer phone calls higher scores. Presumably the owner is more apt to stay on the line with someone he is close to than with a stranger attempting to gain undue access.

Unfortunately, such attacks may still be possible using social engineering (e.g., a stranger may be able to keep the owner on the phone line or get the owner to initiate further communication under false pretenses, such as telling the owner that he has won a contest). Hence, it is important to limit the damage that could be caused by such an attack being successful.

Mitigation 2: Manual verification for high-concern resources. There is currently a steep trade-off between security and convenience depending on whether CARA's classifications are used to directly set a user's social relation (and hence permissions), or whether the classification is only suggested to the owner, who must manually verify that the classification is correct before any permissions are granted. Obviously, requiring manual verification is more secure, and makes most attacks against CARA's classification system a moot point. However, it also reduces the usefulness of CARA.

In our survey results, we found that for some resources, the owner would like to perform some access control, but is relatively unconcerned whether someone gains undue access. This suggests a middle ground where CARA tentatively assigns a social relation to a visitor, but only grants access to such low-concern resources until the social relationship is manually verified by the owner (e.g., a visitor who has been automatically identified as a close friend may then be able to access appliances inside the house, but not be able to unlock the front door).

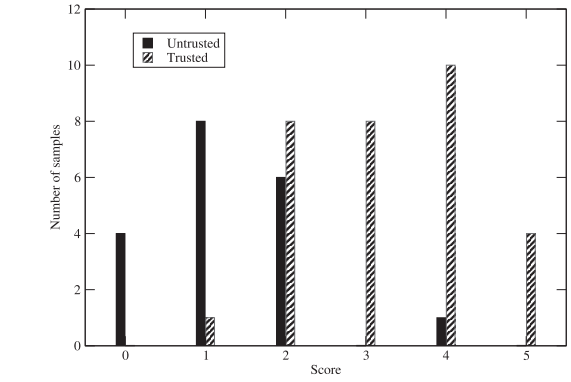

(a)

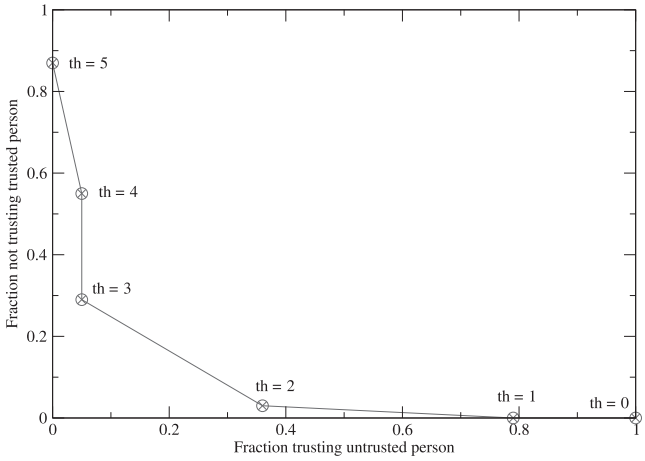

(b)

Fig. 3. Online user study results. The histogram shows the regularity scores for trusted and untrusted people. The ROC plot depicts the error rate for different settings of thresholds for regularity scores. For a given regularity score threshold, the x-axis measures the fraction of untrusted people that would be classified as trusted, and the y-axis measures the fraction of trusted people that would have been classified as untrusted: (a) Histogram of regularity scores and (b) ROC plot.

B. High Value Targets

Certain entities may have unusual communication patterns which make them likely to be misidentified by CARA. Consider, for example, a doctor's office. A CARA user who frequently calls his doctor's office may cause CARA to misidentify the doctor's office as a close friend. From a single CARA system's point of view, this misidentification may seem to be of little consequence, since presumably the doctor's office is unlikely to misuse the resulting credentials. However, if *many* CARA households perform this misidentification, the doctor's office's credentials become a high value target for attackers. An attacker who steals those credentials gains undue access to many homes.

This problem is also mitigated by requiring manual verification for access to high-value targets (as described in subsection VIII-A). In particular, this scenario strongly suggests that CARA ought not to automatically give the ability to enter the owner's home (e.g., unlock exterior doors and windows), nor the ability to access other resources from outside the home.

IX. DISCUSSION

One problem in maintaining access control systems is that they tend to become overpermissive over time. Users who are assigned underpermissive permissions will complain and get it fixed. Users who are assigned overpermissive permissions will not. In real-world systems, this problem is addressed by administrators periodically performing manual audits of who has access to a given resource, and by setting permissions to expire unless manually refreshed [15]. In a home system, manual auditing is likely too much to expect of the owner. The ability to grant temporary permissions, e.g., to a visiting electrician, helps to address this issue but does not entirely solve it. Users are likely to grant indefinite permission to their friends. However, their friends today may not be their friends ten years from now, or even ten weeks from now.

To address this issue, we propose to extend CARA to automatically revoke permissions from users who no longer appear to have as close of a social relationship with the owner (optionally prompting the owner before actually performing the revocation). This can be done simply by having CARA re-evaluate each user given up-to-date social networking information.

A related problem is that the owner may grant a user some elevated permission permanently instead of temporarily, either by mistake or to avoid the trouble of having to refresh expired permissions. After some time, the owner may forget that they granted that user access to the resource, particularly if it is rarely used. Over time, many users may gradually gain ever-inflated permissions. We propose to address this problem by revoking or downgrading (e.g., from unrestricted to logged) access to resources that have not been used for an extended period of time.

X. RELATED WORK

To the best of our knowledge, there has been no previous work that has 1) categorized the complete set of access control policies and classes of visitors specifically for home environments; and 2) proposed a method for automated access policy assignment based on the home owner's social networking data. We discuss here some related work on trust-based access control and policy management for both corporate and home environments.

Many researchers have worked on trust-based security establishment mechanisms. Seigneur *et al.* have developed SECURE framework that has focused on allowing access rights among previously unknown principals to minimize security configuration [16]. Their approach is based on dynamically assessing the trustworthiness of an entity based on three sources of trust: Observation, recommendation, and reputation. Combined with the entity recognition feature for establishing a basis of trust, authors claim that SECURE can provide automated trust establishment. However, adjusting trust based on reputation as described in their paper has some security vulnerabilities; an unauthorized person may be able to gain high trust by stealing a security object that belongs to the home owner and mimicking the owner's biometric information such as his/her voice. Ananthanarayanan *et al.* have proposed an application-level protocol, called SPACE, that enables two devices to establish an automatic ad-hoc secure connection based on the address book entries of two devices [17]. SPACE builds a trust relationship between two devices if they store each other's contact details in their address books without explicit user intervention. However, their system has a flaw that two users can establish a secure connection even though they do not trust each other. CARA minimizes such misconfiguration since CARA evaluates trust between two users based not only on the existence of an entry in each other's address book, but also on the quality and the quantity of the communications between two users.

There has been some work on using portable devices to control access to physical spaces [18]–[21]. Beaufour *et al.* and Zhu *et al.* have considered digital key systems using a mobile device as an alternative to physical key systems to unlock doors [20], [21]. Bauer *et al.* have also used mobile devices as access control tokens for physical space in an office environment [18]. They also conducted a user study in which they derived users' ideal access policies, which included some, such as the *ask for permission* policy, that we discuss in this paper. Further, they showed that these ideal policies could be implemented more accurately and securely with a smartphone-based system than with physical keys [19]. However, their work has focused chiefly on controlling access to a single type of resource (office doors) and only in an office environment.

Some researchers have created tools to assist policy professionals in creating policies. Examples of such tools are the SPARCLE policy workbench, which converts policies in natural language into implemented policy [10]; and the expandable grid, which helps users manipulate implemented policy for file system rules [11]. However, these tools are not built specifically for home environments and also do not target non-expert users.

A number of works have attempted to develop automated or semi-automated means for creating or improving the quality of access control policies. Researchers have developed a number of tools for testing or validating firewall policies against administrator intentions (e.g., [22]–[25]). Bauer *et al.* have proposed using the history of accesses to physical spaces to suggest to users the creation of policies that may be needed in the future [26]. Works on *role mining* (e.g., [27], [28]) use machine-

learning techniques to improve the quality of policies by extracting roles from implemented policy. These roles may then serve as a guide when migrating a legacy system to one that supports role-based access control. All of these works typically targeted a different context and user population than we do.

Several researchers have worked on building ubiquitous home computing systems [29], [30]. Their main focus is to allow household members to configure and manage the introduction and arrangement of new interactive devices and services to meet their own needs, but not on managing access control policies. Argyroudis and O'Mahony have built a system called AETHER, which addresses the establishment of security associations between a set of access control attributes and principals for ubiquitous smart home environments [31]. Although AETHER is one of the most related works in the aspect of providing a foundational architecture for managing security relationships in smart home environments, our work addresses the problem in more detail, such as suggesting a complete set of access control policies and classes of principals. Kostianinen *et al.* have user tested several access control concepts and proposed an access control solution for home networks that imposes minimal burden on the user [32]. However, they have focused on establishing a home network for family members only, and they do not address the automated access of the visitors, which is the core challenge for an efficient and easy-to-use home access control system. Similarly, Marin *et al.* have proposed a home automation middleware for secure management of user and contextual data that gives access to services just to the authorized users and devices [33]. However, their system only considers owners of devices as authorized users and does not cover providing authorization to visitors. Johnson and Stajano have started an initial work on addressing the issues of defining the permissions for guests [34], but this early study does not address technical details.

Based on a user study, Karlson *et al.* explore how security and data privacy concerns affect participants' willingness to share their mobile phones [14]. Their study presents the diversity of guest user categorizations and associated security constraints expressed by the participants. Despite the complexity/impossibility of grouping guests into fixed categorizations, authors argue that 1–3 access settings address much of variations from grouping guests.

Some researchers have studied access control within the home environment. Brush and Inkpen present results from an empirical study of 15 families, and discuss about the degree of shared ownership and use of technologies that families own [35]. Their result suggests that families often share the ownership of technology in the public living space and they trust their family members. As a result, people maintain separate profiles on technologies only to prevent teenagers from accessing computer or to prevent malicious outsiders. Moreover, families who use separate profiles have trouble sharing sharing files/applications with other family members. Also, none of the households use password for personal privacy but rather for personalization. Authors suggest that future devices may need to support both the shared usage and the ability to access personal profile.

XI. CONCLUSION

With our suggestion, Alice and Bob can now feel safe about their McMansion. They do not need to worry about Carol, David, and Evan from manipulating the security device settings, and the number of hours their kids spend in front of TV and the Internet will be restricted as Alice and Bob wish. They can also easily change the control setting for Frank, probably from "minimal control" to "partial control" while reassuring that he would not access their tax files. Similarly, they feel safe that Helen's improper actions are either prohibited or logged for their review. For Irene, Bob can review the log on what she checks at home, and with the level of accessibility, she does not misinterpret that Bob distrusts her. It is not just Alice and Bob who are satisfied; their child Carol also feels safe from George by restricting his access on her photos and diary.

We observe that providing access to home resources to visitors is a challenging research problem, mainly because of the heterogeneity and complexity of home resources, the diversity of visitors, the distrust revelation problem, and the inexperience in security of the home owner. Without sensible mechanisms, visitors could either obtain access to sensitive personal data (in the case of liberal access assignment), or not be able to use the light switch (in the case of restrictive access assignment).

In this paper, we provide an approach to address some of these challenges by assigning visitors access rights from one of four pre-defined control settings, each constructed using one of three proposed policy types. Although our proposed access assignment mechanism based on three dimensions of policies and intuitive access control settings is easy to use, we found that social aspects and scalability issues provide compelling arguments for automation. Our proposal to leverage the call log of cell phones to determine the social relationship between a visitor and a home owner is a promising mechanism. However, further research is needed to identify the ideal mechanism for this purpose.

We anticipate that the research community will embrace this important research challenge to make future home networks at least as secure and as usable as current homes.

REFERENCES

[1] M. Weiser, "The computer for the twenty-first century," *Scientific American*, vol. 265, pp. 94–104, Sept. 1991.

[2] T. H.-J. Kim, L. Bauer, J. Newsome, A. Perrig, and J. Walker, "Challenges in access right assignment for secure home networks," in *Proc. USENIX HotSec*, 2010.

[3] D. Balfanz, D. K. Smetters, P. Stewart, and H. C. Wong, "Talking to strangers: Authentication in ad-hoc wireless networks," in *Proc. NDSS*, 2002.

[4] J. M. McCune, A. Perrig, and M. K. Reiter, "Seeing-is-believing: Using camera phones for human-verifiable authentication," in *Proc. IEEE Symp. Security and Privacy*, 2005.

[5] M. Blaze, J. Feigenbaum, and A. D. Keromytis, "KeyNote: Trust management for public-key infrastructures," in *Proc. Int. Workshop on Security Protocols*, 1999.

[6] M. Blaze, J. Feigenbaum, J. Ioannidis, and A. D. Keromytis, "The KeyNote trust management system," Internet Request for Comment RFC 2704, Internet Engineering Task Force, 1999.

[7] M. Blaze, J. Feigenbaum, and J. Lacy, "Decentralized trust management," in *Proc. IEEE Symp. Research in Security and Privacy*, 1996.

[8] Trusted Computing Group, "Trusted platform module main specification, Part 1: Design principles, Part 2: TPM structures, Part 3: Commands," Version 1.2, Revision 103, 2007.

[9] P. Bergstrom, K. Driscoll, and J. Kimball, "Making home automation communications secure," *Computer*, vol. 34, no. 10, pp. 50–56, 2001.

[10] C. A. Brodie, C.-M. Karat, and J. Karat, "An empirical study of natural language parsing of privacy policy rules using the sparcle policy workbench," in *Proc. The Second Symp. Usable Privacy and Security*, 2006.
[11] R. W. Reeder, L. Bauer, L. F. Cranor, M. K. Reiter, K. Bacon, K. How, and H. Strong, "Expandable grids for visualizing and authoring computer security policies," in *Proc. Conf. Human Factors in Comput. Syst.*, 2008.
[12] J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, and L. F. Cranor, "Crying wolf: An empirical study of SSL warning effectiveness," in *Proc. USENIX Security*, 2009.
[13] M. L. Mazurek, J. Arsenault, J. Breese, N. Gupta, I. Ion, C. Johns, D. Lee, Y. Liang, J. Olsen, B. Salmon, R. Shay, K. Vaniea, L. Bauer, L. F. Cranor, G. R. Ganger, and M. K. Reiter, "Access control for home data sharing: Attitudes, needs, and practices," in *Proc. Int. Conf. Human Factors in Comput. Syst.*, 2010.
[14] A. K. Karlson, A. B. Brush, and S. Schechter, "Can I borrow your phone?: Understanding concerns when sharing mobile phones," in *Proc. Int. Conf. Human Factors in Comput. Syst.*, 2009.
[15] L. Bauer, L. Cranor, R. W. Reeder, M. K. Reiter, and K. Vaniea, "Real life challenges in access-control management," in *Proc. Conf. Human Factors in Comput. Syst.*, 2009.
[16] J. M. Seigneur, C. D. Jensen, S. Farrell, E. Gray, and Y. Chen, "Towards security auto-configuration for smart appliances," in *Proc. The Smart Objects Conf.*, 2003.
[17] G. Ananthanarayanan, R. Venkatesan, P. Naldurg, S. Blagsvedt, and A. Hemakumar, "Space: Secure protocol for address-book based connection establishment," in *Proc. HotNets*, 2006.
[18] L. Bauer, S. Garriss, J. M. McCune, M. K. Reiter, J. Rouse, and P. Rutenbar, "Device-enabled authorization in the grey system," in *Proc. The Int. Conf. Inf. Security*, 2005.
[19] L. Bauer, L. Cranor, R. W. Reeder, M. K. Reiter, and K. Vaniea, "A user study of policy creation in a flexible access-control system," in *Proc. Conf. Human Factors in Comput. Syst.*, 2008.
[20] A. Beaufour and P. Bonnet, "Personal servers as digital keys," in *Proc. IEEE Int. Conf. Pervasive Comput. Commun.*, 2004.
[21] F. Zhu, M. W. Mutka, and L. M. Ni, "The master key: A private authentication approach for pervasive computing environments," in *Proc. IEEE Int. Conf. Pervasive Comput. Commun.*, 2006.
[22] A. Mayer, A. Wool, and E. Ziskind, "Fang: A firewall analysis engine," in *Proc. IEEE Symp. Security and Privacy*, 2000.
[23] S. Hazelhurst, A. Attar, and R. Sinnappan, "Algorithms for improving the dependability of firewall and filter rule lists," in *Proc. Int. Conf. Dependable Systems and Networks*, 2000.
[24] E. S. Al-Shaer and H. H. Hamed, "Firewall policy advisor for anomaly detection and rule editing," in *Proc. Int. Symp. Integr. Network Manage.*, 2003.
[25] F. Le, S. Lee, T. Wong, H. S. Kim, and D. Newcomb, "Minerals: Using data mining to detect router misconfigurations," in *Proc. SIGCOMM Workshop on Mining Network Data*, 2006.
[26] L. Bauer, S. Garriss, and M. K. Reiter, "Detecting and resolving policy misconfigurations in access-control systems," in *Proc. ACM Symp. Access Control Models and Technol.*, 2008.
[27] M. Kuhlmann, D. Shohat, and G. Schimpf, "Role mining–revealing business roles for security administration using data mining technology," in *Proc. ACM SACMAT*, 2003.
[28] J. Schlegelmilch and U. Steffens, "Role mining with ORCA," in *Proc. ACM SACMAT*, 2005.
[29] R. Campbell, J. Al-Muhtadi, P. Naldurg, G. Sampemane, and M. D. Mickunas, "Towards security and privacy for pervasive computing," in *Proc. Theories and Systems, Mext-NSF-JSPS (ISSS)*, 2002.
[30] M. Romn, C. K. Hess, R. Cerqueira, A. Ranganathan, R. H. Campbell, and K. Nahrstedt, "Gaia: A middleware infrastructure to enable active spaces," *IEEE Pervasive Comput.*, vol.1 no. 2, pp. 74–83, 2002.
[31] P. Argyroudis and D. O'Mahony, "Securing communications in the smart home," in *Proc. EUC*, 2004.
[32] K. Kostiainen, O. Rantapuska, S. Moloney, V. Roto, U. Holmstrom, and K. Karvonen, "Usable access control inside home networks," Nokia Research Center, Tech. Rep. NRC-TR-2007-009, 2007.
[33] A. Marin, W. Mueller, R. Schaefer, F. Almenarez, D. Diaz, and M. Ziegler, "Middleware for secure home access and control," in *Proc. IEEE Int. Conf. Pervasive Comput. and Commun. Workshops*, 2007.
[34] M. Johnson and F. Stajano, "Usability of security management: Defining the permissions of guests," in *Proc. The Security Protocols Workshop*, 2006.
[35] A. J. B. Brush and K. M. Inkpen, "Yours, mine and ours? Sharing and use of technology in domestic environments," in *Proc. Ubicomp*, 2007.

Tiffany Hyun-Jin Kim is a Ph.D. candidate in the Electrical and Computer Engineering Department at Carnegie Mellon University. She earned a M.S. degree in Computer Science from Yale University, and a B.A. degree in Computer Science from University of California at Berkeley. Her research covers trust management and usable security and privacy. Her other projects include access control for home networks, security in vehicular networks, and security for mobile and wireless ad-hoc networks.

Lujo Bauer is an Assistant Research Professor in CyLab and the Electrical and Computer Engineering Department at Carnegie Mellon University. He received his B.S. in Computer Science from Yale University and his Ph.D., also in Computer Science, from Princeton University. His research interests cover many aspects of computer security, including distributed access control, policy specification, run-time enforcement, and usable security. He is particularly interested in building usable access-control systems with sound theoretical underpinnings, and generally in narrowing the gap between a formal model and a usable system.

James Newsome is a Systems Scientist in Cylab at Carnegie Mellon University. He received his B.S. degree from the University of Michigan in 2002 and his Ph.D. degree from Carnegie Mellon University in 2008. His primary research interest is in trustworthy computing.

Adrian Perrig earned his Ph.D. degree in Computer Science from Carnegie Mellon University. Currently, he is a Professor in Electrical and Computer Engineering, Engineering and Public Policy, and Computer Science at Carnegie Mellon University. He serves as the technical director for Carnegie Mellon's Cybersecurity Laboratory (CyLab). He is a recipient of the NSF CAREER award in 2004, IBM faculty fellowships in 2004 and 2005, and the Sloan research fellowship in 2006.

Jesse Walker is a Principal Engineer in Intel Corporation's Security Research Lab. He is a co-designer of Skein, one of the finalist algorithms in the SHA-3 competition. He was the first person to identify the security flaws in WEP, the original IEEE 802.11 encryption scheme, and served as technical editor for its replacement, 802.11i. He received a Ph.D. in mathematics from the University of Texas.