

©Copyright 2014

Tamara Denning

Human-Centric Security and Privacy for Emerging Technologies

Tamara Denning

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2014

Reading Committee:

Tadayoshi Kohno, Chair

Batya Friedman

James Fogarty

Program Authorized to Offer Degree:
UW Computer Science & Engineering

UMI Number: 3641517

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI 3641517

Published by ProQuest LLC (2014). Copyright in the Dissertation held by the Author.

Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against unauthorized copying under Title 17, United States Code



ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

University of Washington

Abstract

Human-Centric Security and Privacy for Emerging Technologies

Tamara Denning

Chair of the Supervisory Committee:
Associate Professor Tadayoshi Kohno
Computer Science & Engineering

The creation and adoption of connectivity-, sensor-, and actuator-rich emerging technologies alter the landscape for computer security and privacy. New technologies facilitate novel or amplified kinds of attacks on the financial, physical, and emotional wellbeing of users and people in other, non-user roles. Moreover, the fast rate at which the security landscape changes can often outpace the understanding of users and technologists. My work seeks to enhance people's security and privacy with emerging technologies. In particular, I take a human-centric approach to designing systems for security and privacy, and a human-centric approach to enabling people to achieve better outcomes.

Effective security is not simply a technical challenge, but also a human one. Designing technical systems without considering the humans involved results in suboptimal or unacceptable security solutions. In addition to prioritizing usability, designing good security means designing effective systems that will be embraced by users, fit into the application context, and have minimal negative side effects; this approach requires a deeper understanding of the people in and around a system, their values, and the contexts of technology use.

My thesis work contributes to security and privacy for emerging technologies in two ways: via inductive investigations to support designing security and privacy systems

that respect a broader set of needs and values, and via designing and evaluating a tool to increase security awareness. I present my work in security for implantable cardiac devices, in which I use semi-structured interviews and group workshops to elicit contextual information from two different stakeholder groups: cardiac patients and medical providers. Second, I present my work investigating the perspectives of bystanders on augmented reality devices and lay out potential design axes for privacy-mediating technologies. I conclude by addressing the design, production, distribution, and evaluation of Control-Alt-Hack—a tabletop card game targeted to help disseminate high-level security concepts.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	iv
Chapter 1: Introduction	1
1.1 Research Stance: Beyond Usable Security	2
1.2 Contributions	4
Chapter 2: Implantable Medical Devices: Tailoring the Security Design to the Context	6
2.1 Motivation and Overview	7
2.2 Security System Concepts	13
2.3 Patient Perspective	18
2.4 Medical Provider Perspective	31
2.5 Recommendations for Design	47
2.6 Summary	52
2.7 Acknowledgments	53
Chapter 3: Augmented Reality and Bystander Privacy: Exploring Privacy Attitudes With an Emerging Technology	54
3.1 Motivation and Overview	55
3.2 In-Situ Interviews with Bystanders	59
3.3 Design Axes for Privacy-Mediating Technologies	73
3.4 Summary	78
3.5 Acknowledgments	79
Chapter 4: Control-Alt-Hack: A Tabletop Card Game for Computer Security Awareness and Education	80

4.1	Overview	81
4.2	Game Design	82
4.3	Game Evaluation	94
4.4	Summary	113
4.5	Acknowledgments	114
Chapter 5: Conclusion		115
Bibliography		118
Appendix A: Extended Category Details for Medical Provider Workshop Study		129
A.1	Initial Perspectives	129
A.2	Metaphors	131
Appendix B: Additional Control-Alt-Hack Material		133
B.1	Card Topics and Research Papers	133
B.2	Educator Survey Contents	134
B.3	Educator Survey Coding Disagreements	136
B.4	Control-Alt-Hack-themed Assignments	138

LIST OF FIGURES

Figure Number	Page
2.1	Photos of system concept mockups used in the patient interviews. . . 19
2.2	Photos of system concepts used in the medical provider workshops. . . 35
3.1	Photo of AR mockup. 58
3.2	Design axes for privacy-mediating technologies. 73
4.1	Photo of Control-Alt-Hack. 89
4.2	Character art from Control-Alt-Hack. 90
4.3	Sample cards from Control-Alt-Hack. 91

LIST OF TABLES

Table Number	Page
2.1 Security system concepts, relevant properties, and related research.	13
2.2 Participants' evaluations of different system mockups.	23
2.3 Medical provider workshop participant distribution.	32
2.4 System concept evaluations by medical providers.	40
4.1 Educator participants reporting on classroom usage	95
4.2 Educator participants reporting on non-classroom usage	96
4.3 Classroom-based educator survey analysis results	98
4.4 Non-classroom-based educator survey analysis results	101
4.5 Overall educator evaluations of Control-Alt-Hack	106
4.6 Mapping of design goals to user responses	109
A.1 Initial Perspectives from medical provider workshops	131
A.2 Metaphors from medical provider workshops	132
B.1 Example technical research and corresponding card topics	133

ACKNOWLEDGMENTS

I thank everyone who has given advice and guidance and has helped make this work possible. In particular, I (alphabetically) thank all the people with whom I have collaborated: Adrienne Andrew, Alan Borning, Gaetano Borriello, Kevin Bowers, Ryan Calo, Rohit Chaudhri, Zakariya Dehlawi, Marten van Dijk, Glen Duncan, Batya Friedman, Kevin Fu, Brian T. Gill, Mark Gondree, Carl Hartung, Ari Juels, Tadayoshi Kohno, Karl Koscher, Daniel B. Kramer, Adam Lerner, Jonathan Lester, Henry M. Levy, William H. Maisel, Yoky Matsuoka, Cynthia Matuszek, Bryce C. Newell, Zachary N.J. Peterson, Matthew R. Reynolds, Adam Shostack, and Joshua R. Smith.

I thank members of the University of Washington Computer Security and Privacy Research Lab.

Other people helped immensely; particular chapters provide more complete acknowledgments as they pertain to specific work.

I was supported during this work by an Achievement Rewards for College Scientists (ARCS) Fellowship, the Marilyn Fries Endowed Regental Fellowship, a Microsoft Research Womens Scholarship, and an Intel Ph.D. Fellowship. Funding for particular work is specified at the end of chapters.

DEDICATION

to the University of Washington

Computer Security and Privacy Research Lab—my academic family.

Chapter 1

INTRODUCTION

Computer security and privacy are a critical component of the successful design, deployment, and usage of technology systems: failures in a technology's security or mismatches in privacy expectations result in a variety of financial, physical, and emotional harms. In turn, this can erode trust in—and adoption of—new technologies.

To further complicate matters, the landscape of technologies is always changing. Emerging technologies are increasingly interconnected, incorporate a wider variety of sensors and actuators to interface with the physical world, and are incorporated more broadly and deeply into our lives. While these properties provide utility, they also have implications for security and privacy. Adversaries are able to execute attacks more easily and at scale, and people can be impacted in novel ways.

In many cases, these emerging technologies require new approaches to security and privacy; differences in attack surfaces, differences in interfaces, and differences in the ability to affect human assets all contribute to creating different security scenarios. To compound the problem, the pace of change overwhelms people's familiarity with the security and privacy landscape.

My dissertation work seeks to enhance people's security and privacy with emerging technologies—technologies beyond the desktop or the laptop with characteristics such as being wearable, implanted, sensor-rich, or cyber-physical. Additionally, I take a human-centric approach to designing systems for security and privacy that embody a more holistic perspective, and I take a human-centric approach to enabling people to achieve better security and privacy outcomes.

In my dissertation, I begin by addressing an important class of current technologies

which is both implanted and cyber-physical: implantable medical devices. I investigate how mixed methods—that is, both quantitative and qualitative methods—can be used to capture domain knowledge and generate design recommendations that inform the design of better security: systems that better satisfy the needs, preferences, and constraints of multiple stakeholders surrounding the technology. Next, I investigate how mixed methods can be used to gather people’s perspectives on the privacy implications of a technology that is wearable and sensor-rich, but not yet prominent: augmented reality. This work contributes to understanding people’s reception of the technology and the specific concerns that they have regarding its capabilities; as with my work on implantable medical devices, these results inform the design of systems that are more respectful of both users and bystanders and that minimize their negative side effects.

A common problem that I encountered throughout my work is that people are insufficiently aware of the potential impacts of security failures with emerging technologies. This problem can be addressed in different ways, including by designing more intuitive interaction experiences or by directly delivering educational information; I wished to investigate how alternative tools could be used to deliver small amounts of security information in a more implicit manner. I designed, produced, and distributed Control-Alt-Hack: a tabletop card game targeted to impart high-level security concepts via a more engaging activity. Subsequently, I evaluated the usage of the game in educational contexts.

Together, the elements of my dissertation all contribute towards creating better security and privacy outcomes with emerging technologies via studies and tools that have a human-centric focus.

1.1 Research Stance: Beyond Usable Security

Effective security is not simply a technical challenge, but also a human one. Designing technical systems without considering the humans involved can result in suboptimal

or unacceptable security solutions. From a usability standpoint, users might unintentionally misconfigure or misuse a system; if the system is sufficiently frustrating or incomprehensible, they might intentionally deactivate it (e.g., [20, 93]).

Beyond usability, however, designing good security means designing effective systems that will be embraced by users, fit into the application context, and have a minimum of negative side effects; this more holistic approach requires a deeper understanding of the people in and around a system, their values, and the contexts of technology use. This understanding is especially valuable with emerging technologies, which may incorporate new capabilities, be used in new contexts, or be less well understood. In order for the security community to design effective and realistic security systems, it is critical that human beings and human-centered methods be incorporated into the research process.

Prior research has explored users' internal mental models in order to help security researchers design systems that react in alignment with users' expectations; this kind of work also helps security researchers understand the reasons behind incorrect user behaviors (e.g., [34, 90]). Other work examines how to design for and evaluate the performance of usable security and access control systems (e.g., [21, 75]). While this body of research is valuable, the studies are often focused solely on users, and often give less consideration to other important stakeholders or the larger ecosystem in which a technology and its security system are likely to be deployed.

My research approaches design for security and privacy with a focus on:

- **Considering Multiple Stakeholder Roles.** Who can affect, or is affected by, the technology in question? This may include people who are not direct users of the system. What does this mean in terms of system design? In what roles do people interact with the technology, and in what roles do they affect—or are they affected by—the system?
- **Considering Nuances of Specific Usage Contexts.** What is the specific

context in which the technology is used? How does it affect—or how is it affected by—the desired or functional security and privacy properties? What are the implications for system design?

- **Considering Non-Security Factors.** When considering multiple stakeholder roles and the specific usage contexts of a technology, what other factors—besides security or privacy—are important to people? How do different security and privacy system designs impact these factors?
- **Communicating Technical Concepts in Comprehensible Ways.** How can security researchers and designers communicate about technical security concepts with people in ways that are understandable and effective? How can we evaluate the performance of such techniques—particularly when the desired outcomes are more nebulous in nature?

1.2 Contributions

My thesis work makes the following contributions to security and privacy for emerging technologies:

Domain Knowledge to Support Designing for Security and Privacy. My work improves the foundation of knowledge for security and privacy of emerging technologies in two specific domains. Chapter 2 details my work with security for implantable cardiac devices and Chapter 3 lays out my work with augmented reality and privacy.

Emphasizing More Holistic Security Design. Throughout my work, I emphasize taking a more holistic approach to designing for security and privacy. This approach incorporates the consideration of multiple stakeholders in different roles, the consideration of the nuances of specific usage contexts, and the consideration of how security and privacy interacts with other needs and values.

Tools to Increase Security Awareness. Chapter 4 of this work deals with creation of a non-standard tool to help disseminate high-level security information. Specifically, I address the design, production, distribution, and evaluation of Control-Alt-Hack—a tabletop card game targeted to deliver an awareness of the breadth of technologies impacted by computer security, the creativity of attackers, and the various ways that people can be harmed by security and privacy failures.

Chapter 2

IMPLANTABLE MEDICAL DEVICES: TAILORING THE SECURITY DESIGN TO THE CONTEXT

The work in this chapter of my dissertation deals with designing better security systems for a specific domain: implantable cardiac devices. These devices are a type of emerging technology that is implanted, cyber-physical, deals with private health information, and has direct health impacts. Designing better security for implantable medical devices is timely, since they are incorporating increasing connectivity. The medical setting is also a particularly pertinent domain for the application of elements of my research stance: considering multiple stakeholder roles, considering nuances of specific usage contexts, and considering other factors aside from security.

Section 2.1 provides background information on implantable medical devices and security for implantable medical devices, as well as background on value sensitive design, which framed this work. Additionally, this section describes the security system concepts that are used in the studies described in the subsequent sections. Section 2.3 deals with investigating the perspectives of patients who have implanted cardiac devices. The methodology used in this section is semi-structured interviews. Section 2.4 details the study I performed to capture domain expertise from medical providers. The methodology used in this section is group workshops. Section 2.5 synthesizes the results from Sections 2.3 and 2.4 to provide recommendations for the design of security systems for implantable cardiac devices. Section 2.6 summarizes this chapter in the context of my dissertation.

Part of the material in this chapter first appeared in [23] and [25]. Collaborators

on the Cloaker (fail-open wristband) work included Kevin Fu and Tadayoshi Kohno. Collaborators on the patient study included (in alphabetical order): Alan Borning, Batya Friedman, Brian T. Gill, Tadayoshi Kohno, and William H. Maisel. Collaborators on the provider study included (in alphabetical order): Batya Friedman, Brian T. Gill, Tadayoshi Kohno, Daniel B. Kramer, and Matthew R. Reynolds.

2.1 Motivation and Overview

Implantable medical devices (IMDs), such as pacemakers and implantable cardioverter-defibrillators (ICDs), are electronic devices designed to treat abnormal physiological conditions within the body. Millions of people worldwide depend upon IMDs for life-critical functions, and we can only expect this number to grow in the future.

Security for these devices is critical: they handle protected health information (PHI), they are expensive (and non-trivial) to replace, and they are specifically designed to cause physiological changes in people's bodies. The computer security research community has begun to investigate technical security designs suited to the computational requirements and restrictions of these embedded devices. Developing strong technical defenses is, however, only part of the solution. There is a fundamental gap between developing technical mechanisms that *could* protect the security of future wireless medical devices if deployed and developing security defenses that *will* be accepted (even welcomed) by patients, doctors, and other stakeholders. Moreover, any technical security system must be suited to its ecosystem and cause a minimum of negative side effects. This is particularly true of a domain like implantable medical devices, where the attacks have yet to manifest and where the costs of poorly designed security systems translate into very real repercussions on time, effort, and inaccessibility.

This research seeks to help bridge this gap between technical systems and effective deployments by initiating an analysis of how potential security systems for IMDs interact with the needs and restrictions of two critical stakeholder groups: cardiac

patients and medical providers. In order to frame these studies, my collaborators and I used value sensitive design (e.g., [33, 35, 63, 64])—a framework from human-computer interaction that emphasizes prioritizing a wider array of human values and investigating the impacts of technologies from multiple stakeholder perspectives. The contributions of these studies are as follows:

Domain Results. The results from these studies: (1) provide feedback on current technical directions of security for implantable medical devices from important stakeholders, thereby “closing the loop” for system design; (2) provide some domain context and terminology, which security researchers can use to orient themselves and communicate with experts in the medical domain; and (3) provide concrete design recommendations for security researchers working in the implantable medical device space.

Methodology. The methodologies used in these studies are selected and adapted from the human-computer interaction literature in order to elicit pertinent information from relevant stakeholders. While much of the gathered data might be considered common knowledge by those with experience in the area, my collaborators and I refine and utilize procedures with which to gather this information in a structured manner. The stakeholder characteristics and the investigative goals for the patient and provider studies were different, and therefore necessitated different methodological formats. While my collaborators and I have conducted this research in the context of implantable medical devices, the workshop technique could be used to explore other domains of emerging technologies such as automobiles, augmented reality, or 3D printing.

2.1.1 Background on Implantable Cardiac Devices

Current electronic implantable medical devices (IMDs) prevent or treat conditions ranging from heart failure to Parkinson’s Disease. My thesis work focuses on patients with implanted cardiac devices: pacemakers and implantable cardioverter-

defibrillators (ICDs). For some patients, pacemakers might improve their lives by restoring more normal heart rhythms, while in other patients their device might be life-sustaining. Patients with ICDs depend upon their devices to treat potentially fatal heart rhythms.

Implantable cardiac devices store information such as the patient's name and records of irregular heart rhythms that occurred since the last checkup. Patients visit cardiology clinics periodically so that medical staff can download information about these episodes and adjust settings on the patient's device. Current-generation cardiac devices have the ability to communicate wirelessly with external equipment from distances up to 5 meters away via a dedicated FCC band for medical devices. There are numerous reasons for making IMDs wireless. For example, wireless IMDs can be configured ("programmed") by doctors in the operating room from farther away, which avoids the need to bring programming equipment into the sterile operating area. Wireless technology also allows the IMD to send alerts to a home monitoring station—which can then send a report to the patient's physician for analysis—without causing interruption to the patient's activities (including sleep). Unfortunately, incorporating a new wireless interface for these devices also increases the communication surface by which they can be attacked.

2.1.2 Security & Implantable Cardiac Devices: Attacks and Defenses

In 2008, Halperin et al. demonstrated that an implantable cardiac device with centimeters-range wireless communications capabilities can be wirelessly compromised by a (nearby) unauthorized party [45]. The authors showed that someone using low-cost, home-made equipment (i.e., a software-defined radio) could communicate with the ICD to learn private information about the patient, issue large shocks, deactivate potentially life-saving therapies, and induce dangerous heart rhythms. Subsequently, Gollakota et al. demonstrated that an implantable cardiac device with meters-range wireless capabilities is also wirelessly compromisable [39].

Since then, the computer security research community has begun to investigate new technical security mechanisms for warding off the potential security risks to future wireless IMDs. In 2008, Halperin et al. outlined challenges for securing wireless IMDs and possible directions for improving future security [44]. One key challenge they identified is the need to balance *security* (blocking inappropriate access) while also providing some guarantee that *safety* can be ensured in an emergency (facilitating appropriate situational access). To illustrate this point, consider a security system in which the IMD only grants wireless access to individuals who know a password, such as the implanting physician or the follow-up cardiologist. While such a system does improve security and can prevent unauthorized access by random individuals, this system also directly and negatively impacts safety: emergency personnel will not be able to read or change settings on the device without first contacting the patient's cardiologist, who might be unreachable.

There have been numerous early-stage technical proposals designed to help improve security while still facilitating unplanned medical access. One proposed direction requires the patient to wear a wristband that protects the security of the IMD when worn, but that can be removed for emergency access [25, 39, 96]. Another direction requires body modifications, such as RFID implants or tattoos with visible or UV-visible ink [23, 76]. Yet another direction requires the doctor to place something on or near the patient in order to activate longer-range wireless capabilities (e.g., [18, 72, 89, 98]), taking advantage of cryptographic distance-bounding, intra-body signaling, or physiologically-derived keys. Drawing from past work, an IMD could also potentially use automated techniques to detect emergency situations and decrease security requirements (e.g., not require a password) if the patient is in a state of medical emergency [42]. Another, more traditional approach might be to issue temporary or permanent access passwords via a centralized entity, such as a manufacturer-maintained database.

Although less related to this work on implantable cardiac devices, there has been

significant work focused on security and privacy for personal medical sensors and networked medical devices (e.g., [71]); see Avancha, et al. for a survey [8]. Many of the efforts in this space also have potential applicability to implantable medical devices. For example, the Amulet system [83]—which requires the user to wear an external device—has many overlapping elements with other defenses for IMDs [25, 39, 96]. Additionally, any effort to improve key establishment for body-area networks (e.g., [18, 89]), can be used to help improve the security of key establishment systems for IMDs.

Recently, Rushanan et al. published a systematization of knowledge paper on the topic of computer security for implantable medical devices and body area networks (BANs) [74]. The paper provides a thorough examination of both attack and defense work dealing with these classes of technologies in the computer security community. The authors categorize defense directions as falling into four different trends: Biometric and Physiological Values (e.g., ECG or IPI), Out-of-Band (e.g., tattoos), Distance Bounding (e.g., intrabody signaling or cryptographic), and External Devices (e.g., fail-open wristband). Four additional categories—Wireless Attacks, Software/Malware, Anomaly Detection, and Emerging Threats—are used to classify other research trends in the area.

2.1.3 Background on Value Sensitive Design

Computer security and access control systems are frequently discussed in the context of values like security, privacy, and convenience. These systems, however, also affect and are affected by other important human values like trust, physical welfare, autonomy, or human dignity. In this research my collaborators and I drew on established methods from value sensitive design (e.g., [33, 35, 63, 64]) to frame the study design and data analyses.

Value sensitive design was developed in human-computer interaction, and has since been used in civil engineering, information management, human-robotic interaction,

and ubiquitous computing. Value sensitive design has also been used to analyze informed consent for cookies in web browser security, leading to recommendations for browser redesign to better support informed consent and a proof-of-concept redesign in the form of a plug-in “cookie watcher” for the Mozilla browser [33, 64]. Another security-focused study investigated users’ mental models for web browser security, suggesting that elements in the user interface (e.g., the open or closed padlock) were inadvertently leading some users to construct incorrect mental models for a secure connection [34]. More recently, value sensitive design methods have been applied in research about safety and security for mobile phone parenting technologies for teens [22] and home technologies [26].

In this research, my collaborators and I drew explicitly on two value sensitive design methods: direct and indirect stakeholder analyses, and value dams and flows.

Direct and Indirect Stakeholder Analyses. In examining the ecosystem surrounding security for implantable cardiac devices, an important question is what roles should be represented among study participants. Value sensitive design stresses the consideration of multiple stakeholder groups: the direct stakeholders who will interact with the technology (e.g., cardiologists and, to some extent, the patients); and indirect stakeholders who—while they do not directly interact with the technology—can affect and be affected by the technology (e.g., emergency room staff).

Value Dams and Flows. Given a wide range of possible technical security solutions, it is not always obvious how to choose which system to pursue. Value dams and flows is a technique for identifying reasonable, value-sensitive design options from among a large set of possible designs or technical features (e.g., [22, 63]). First, options that a threshold percentage of stakeholders strongly object to are removed from the list of viable solutions (value dams); then, from the remaining options, those that are favored by many stakeholders are selected as good candidates for solutions (value flows). In this research my collaborators and I use the value dams and flows method to help identify viable security designs for implantable cardiac devices.

Security System Property	System Concepts						
	I	II	III	IV	V	VI	VII
Requires patient to wear something [25, 39, 96]	■				■		
Requires modification to the patient's body [76]			■	■			
Requires patient maintenance [25, 39, 96]					■		
Visible on the patient [25, 39, 96]	■		■	■	■		
Depends on centralized infrastructure		■					
Requires specialized equipment [72, 76, 89, 98]			■	■		■	
Requires proximity to the patient [18, 72, 89, 98]	■		■	■	■	■	
Has a manual override [25, 39, 96]					■	■	
Security decisions are automated [42]							■

Table 2.1: Relevant properties of the security system concepts presented to providers, by system concept. Dark cells indicate a property represented by a system. Lighter cells indicate a property represented in some situations or by some interpretations.

2.2 Security System Concepts

In both my work with cardiac patients and with medical providers, I drew upon a common set of security system concepts. These systems are not complete or perfect from a security (access control: false positive), safety (access control: false negative), or usability standpoint. My collaborators and I presented these systems for feedback because: (a) they are representative of some of the security solutions that have been previously proposed by the security research community; (b) they represent a variety of relevant system properties; and (c) the discussion of a specific system can serve to ground an otherwise abstract discussion. Table 2.1 provides a summary of some of the systems' relevant properties and correlates these properties with previous security research. The security concepts are as follows:

- (I) **Medical Alert Bracelet with Password.** Medical alert bracelets are accessories that are worn by some patients in order to inform emergency medical staff of their diagnoses in the case where the patient is unconscious. Since medical

alert bracelets are less likely to be lost in an emergency than materials carried by patients—such as informational cards in their wallets—they are a safer way to convey the IMD passphrase to medical staff. In this security system concept, access to the IMD is protected by a passphrase engraved on the back of a medical alert bracelet that is worn at all times by the patient. This system reasonably satisfies the security property (preventing unauthorized access), since the passphrase can be protected from casual observation. However, the system does not fully satisfy the safety property (access in case of medical emergencies) even if the patient wears the bracelet at all times: the bracelet can still be lost or damaged in an accident, which would render the IMD inaccessible to emergency medical staff.

We included this system in the studies—despite its technical drawbacks—for several reasons: some patients may already wear medical alert bracelets, and people are generally familiar with the concept of a passphrase; the medical alert bracelet provides a contrast to the password-based tattoos (below) since it shares some of their properties but not others; and the medical alert bracelet provides a contrast to the wristband systems, which also require wearing something on the wrist (but operate by a different mechanism and require additional maintenance).

(II) **Centralized Database.** Providers and patients are already familiar with the idea of electronic medical records (EMRs). In this system concept, a centralized database system that medical centers can access could be used to obtain a temporary password which, in turn, could be used to access the implanted cardiac device.

(III) **Tattoo of a Password.** Tattoos have been used throughout history as artistic and cultural forms of expression. They are also an effective way to permanently

carry information. In this security system, access to the IMD would be protected by a passphrase that is encoded as a 2D barcode and tattooed onto the patient's skin. There are two advantages to this system over the medical alert bracelet system: patients cannot forget or lose their passphrase, and they do not have to wear anything on their wrists. This system does not completely satisfy the safety property, since the tattoo could be damaged and rendered unscannable in an accident. Additionally, this system makes it more difficult to revoke or reissue passphrases. Using a tattoo as a security system for medical devices also touches on patient views and values such as self-image and freedom from unwanted historical associations.

My collaborators and I had some reservations about including a tattoo-based patient identifier in these studies, especially considering potential associations with tattooing of prisoners in concentration camps during World War II; however, when some of us discussed the solution space at security conferences, we frequently heard the suggestion that a way to solve the problem of losing a carried passphrase in an accident is to tattoo the passphrase onto the patient's skin. My collaborators and I hypothesized that this system, while somewhat satisfactory from a technical perspective, would not be satisfactory from other perspectives. We included tattoos in these studies in order to confirm or counter our hypothesis.

- (IV) **UV-Visible Tattoo of a Password.** In addition to regular tattoos with black or colored inks, it is now possible to get specialty tattoos that are only visible under ultraviolet (UV) lights. In this system design, access to the IMD is protected by a passphrase that is encoded as a 2D barcode and tattooed onto the patient's skin using an ink that is only visible under a UV light source. This system has an advantage over the visible tattoo system because it cannot be seen under normal conditions, and therefore does not affect the patient's

appearance. This system was included in the studies because it provides a useful contrast to the previous system: it partially decouples a patient’s tattoo from its cultural and historical associations by rendering it non-visible under normal circumstances.

- (V) **Fail-Open Wristband.** I previously performed a high-level investigation of the protocol design space for a potential security system for wireless IMDs. The system design consists of an external computational unit that controls access to the IMD [25]. We called this system a Cloaker because its presence causes the IMD to be “invisible” to all unauthorized queries; however, my collaborators and I chose to call this the “wristband system” in the following studies to avoid using suggestive terminology. The system concept has been expanded upon in subsequent research [39, 96].

When the wristband is present and worn by the patient, the IMD only acts on commands sent by authorized entities. Generally speaking, this is implemented by encrypting all communications to and from the IMD and checking them for authenticity and integrity. When the wristband is removed, the system changes its access status to allow any programmer¹ to access the IMD. In this way emergency medical staff can access a patient’s IMD even if the wristband is lost or destroyed in an accident. This is in contrast to carrying an access passphrase on a medical alert bracelet or card, since emergency access with a traditional bracelet would not be possible if that bracelet or card were forgotten, lost, or destroyed.

- (VI) **Proximity-Bootstrapped Equipment.** This security system consists of an external device that is used by medical staff. When placed in contact with

¹In the context of implantable cardiac devices, the term *programmer* is used to refer to the physical piece of equipment that is used to read and change IMD settings.

the patient, the device negotiates a temporary key with the IMD through the patient's body—for example, via physiological keying [18, 89] or intra-body signaling [44, 98]—thereby gaining permission to access the IMD. As an alternative, this device could use proximity-based access, as proposed by Rasmussen et al. [72]. This security system was included because it does not involve patient participation, uses an external device carried by medical staff, and because a proximity-bootstrapped device is analogous to devices currently used by medical professionals to communicate with IMDs via short- and medium-range wireless.

- (VII) **Criticality-Aware IMD.** This security system concept is a behavior that is built into a pacemaker or ICD, and therefore represents a system that requires no additional patient body modifications, patient behavior changes, or external equipment. The criticality-aware IMD detects indicators such as the patient's location, whether or not the patient is standing up (some pacemakers already incorporate accelerometers), and heart rhythms to determine whether or not the patient is in probable medical distress. If this data suggests that the patient is experiencing a medical emergency, the IMD changes its access policy for the duration of the emergency so that all programmers are authorized to issue commands. This behavior is intended to help ensure that medical staff can access a patient's device in an emergency situation. This concept of a criticality-aware system is similar to Gupta et al.'s work on criticality-aware access for pervasive applications [42] and was proposed for use with IMDs in Halperin et al. [44]. This security system addresses both the security and the safety goals, assuming that it has a low incidence of false negatives (safety failures) and false positives (security failures).

2.3 Patient Perspective

My collaborators and I used the value sensitive design (VSD) framework and semi-structured interviews to explore the values and experiences of patients with implanted cardiac devices and solicit their feedback on a range of security approaches. We developed this interview instrument based on a conceptual investigation of the problem space and informed by our experience as security researchers, as HCI researchers and social scientists, and as a cardiologist. In the course of the interview, I presented participants with 8 different security systems that embody different approaches to security for IMDs. Not all of these systems are well-developed or even desirable as solutions for IMD security; instead, we solicited feedback on them in order to investigate the ways in which participants' values and priorities interact with the security systems' properties. The study data results in 11 criteria that we advise researchers and designers to address in order to make their security systems more acceptable to patients with IMDs. By studying how patients' views and values interact with security systems for wireless implantable medical devices, this work helps serve as a foundation for informing the design of future IMD systems that not only have desirable technical security properties, but that also address the values and needs of patients.

2.3.1 Study Design

System Concepts and Mockups

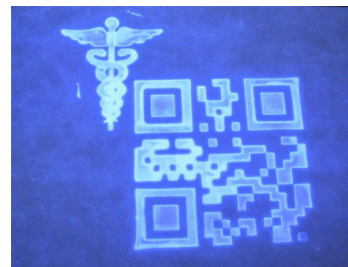
My collaborators and I chose to use mockups of the systems during the interview to help participants envision the future systems [29]. While using specific objects can—and in some interviews did—cause participants to react to the particular appearance of the object rather than the general system properties, we found that the mockups helped make the systems concrete for participants. Since we were presenting 8 different systems, the mockups also functioned as memory triggers to help participants



(a) The back of the medical alert bracelet mockup (System I).



(b) The tattoo mockup (System III).



(c) The UV-visible tattoo mockup (System IV).



(d) The wristband mockups (Systems 5, 5', 5'').



(e) The proximity-bootstrapped equipment mockup (System VI).



(f) The criticality-aware IMD mockup (System VII).

Figure 2.1: Photos of system concept mockups used in the patient interviews.

keep track of the different systems. The system concepts chosen for inclusion in the interview are listed below, along with the category to which they were assigned (for logistical purposes, such as section randomization during the interviews). System Concept Descriptions are given in Section 2.2. Photos of the system mockups used during the interviews are shown in Figure 2.1.

- **Passwords and Additional Patient Body Modifications**

- **System I: Medical Alert Bracelet with Password.** A description of this system (see Figure 2.1a) is given in Section 2.2.
- **System III: Tattoo of a Password.** A description of this system (see Figure 2.1b) is given in Section 2.2.
- **System IV: UV-Visible Tattoo of a Password.** A description of this system (see Figure 2.1c) is given in Section 2.2.

- **Patient Behavior Change**

- **System V: Fail-Open Wristband.** A description of this system (see Figure 2.1d) is given in Section 2.2.
- **System V': Fail-Open/Safety Wristband.** This security system (see Figure 2.1d), as with the previous system, restricts access when it is present and allows open access when it is absent. In contrast to the previous version, however, this system has additional features. It sounds an alarm when a patient enters an environment with a strong magnet, since strong magnets can affect the IMD's operation. The wristband also dials 911 when it detects that the patient is experiencing a cardiac emergency. The purpose of including this system in the interview was to contrast the participant's reactions to the regular wristband system with the participant's reactions to that same system when it offers additional safety benefits.

- **System V'': Fail-Open Wristband with Patient-Specified Functionality.** This security system (see Figure 2.1d), as with the previous two, restricts access when the wristband is present and allows open access when it is absent. This system does not have the safety features of the previous (emergency and warning) system, but I invited participants to specify some additional functionality that they might find useful. In the interview we offer the examples of a watch, a pedometer, and a heart-rate monitor. This version of the wristband system is included in the interview so that we can investigate whether there is some other functionality that might entice patients to wear and maintain a security system.

- **Passive with Respect to the Patient**

- **System VI: Proximity-Bootstrapped Equipment.** A description of this system (see Figure 2.1e) is given in Section 2.2.
- **System VII: Criticality-Aware IMD.** A description of this system (see Figure 2.1f) is given in Section 2.2.

Recruitment and Format

Patients were recruited for the study at the pacemaker and ICD clinic at a large urban hospital on the east coast of the United States after approval of the research protocol by the hospital's Institutional Review Board and the Institutional Review Board at the University of Washington. Patients with clinic appointments during the study window were asked to participate in the study. Participating patients were compensated for their expenses (parking).

We refined the interview protocol by conducting three pilot interviews with cardiac patients. A total of 17 interviews were conducted. Due to incompleteness ($N = 2$) and an irregularity in the questions ($N = 2$) 4 interviews were dropped. The data

for the remaining 13 patients is reported here. Of the 13 patients, 9 had pacemakers and 4 had ICDs. The patients' average age was 67.9 (median = 69, range = 41–80). Our participant population was, on average, on their second implanted cardiac device (median = 1, range = 1–3+) and had lived with a device for 7.8 years (median = 6, range = 0–18).

The semi-structured interview protocol contained a combination of yes/no, multiple choice, and open-ended questions. The Mockup Evaluation portion of the interview presents participants with all 8 security systems in the Password and Body Modification, Patient Behavior Change (Wristband), and Patient-Passive categories. The category order was randomized across participants.

The Mockup Evaluation is divided into two parts. In the Mockup Properties portion, the interviewer explains each system and solicits positive and negative feedback about that system. In the Comparative Mockup Evaluation portion, once all systems have been presented, participants are asked to identify: (1) the systems that they liked; (2) the systems that they disliked; and (3) the system or systems that they would choose to use, if they were asked to use a security system in the future, though some people chose none for (3). The interviews were audio-recorded and later transcribed.

Evaluative (quantitative) responses from the interview were coded in the following process: (1) the primary coder developed a coding scheme for each quantitative question based on the possible answers; (2) the primary coder coded the evaluative responses for all interviews; (3) the reliability coder coded the evaluative responses for all interviews; and (4) Cohen's kappa was computed for the results. The overall value of Cohen's kappa for the quantitative responses reported for this study is 0.75.

2.3.2 Results: Quantitative Evaluations

This section gives the results from the evaluative portion of the interview. The results give some indication of participants' relative preferences for the presented system

Mockup System	Liked (N= 11)	Disliked (N= 11)	Would Choose (N= 11)
I. Medical alert bracelet	0%	27%	0%
III. Visible tattoo	9%	55%	9%
IV. UV-visible tattoo	18%	27%	18%
V. Fail-open wristband	0%	36%	0%
V'. Fail-open/Safety	45%	27%	27%
V''. Fail-open/Patient-specified	0%	36%	9%
VI. Proximity-bootstrapping	27%	0%	27%
VII. Criticality-aware IMD	27%	18%	27%

Table 2.2: Participants' evaluations of different system mockups.

concepts; however, these results should not be overemphasized. The methodology and sample size of the study are geared towards an inductive investigation of the design space and participants' reactions to system concept properties, rather than a deductive ranking of system concepts.

If participants (1) liked the system and did not (2) dislike the system, the combined code was "like." If they did not (1) like the system and (2) disliked the system, the combined code was "dislike." If they did not (1) like the system and did not (2) dislike the system, the combined code was "neither." If they *both* (1) liked the system and (2) disliked the system, the combined code was "neither." If either response was uncodable, the combined code was also uncodable. For this portion of the interview, responses from two participants were uncodable. Thus, $N = 11$ for these analyses. The results of these codings are shown in Table 2.2.

Dams and Flows Analysis

The Least Disliked: Proximity-Bootstrapped Equipment. In the absence of a consensus on a liked system, my collaborators and I use the technique of VSD dams and flows as the inspiration to isolate multiple systems that would achieve

more complete “coverage” for the participants. We focused first on “dams,” since it is particularly important at a minimum that every patient have a choice of a system that they do not actively dislike. As shown in Table 2.2, the proximity-bootstrapped equipment concept is not the most liked of the systems, but the fact that no participants disliked the concept makes it a good candidate for part of a security solution for implantable cardiac devices. This security solution, however, is not optimal: only 27% (3 out of 11) of the participants liked it and would choose to use it.

The Most Liked: Fail-Open/Safety Wristband. Once my collaborators and I had a system that none of the participants in our sample disliked, we shifted our attention to “flows,” seeking to select additional system concepts to create a portfolio of security systems that give a higher percentage of participants an option that they would like or choose to use. From inspection of the data (see Table 2.2), the most natural choice appears to be the fail-open/safety wristband concept, since it has high percentages in both categories. Another possible candidate might be the criticality-aware IMD. However, if we choose two concepts as system options, we should avoid concepts which tend to be liked, disliked, and/or selected by the same participants, since choosing both systems would be redundant. One way to measure this is to consider the correlation between the like/dislike scores for pairs of concepts and avoid selecting pairs for which scores are highly positively correlated. The like/dislike ratings for the criticality-aware IMD are positively correlated with the proximity-bootstrapped equipment (using Kendall’s tau-b, a non-parametric correlation coefficient, $\tau = 0.510$), so this concept would not be a useful system to add to the solution portfolio. On the other hand, the like/dislike ratings for the proximity-bootstrapped equipment and the fail-open/safety wristband are essentially uncorrelated ($\tau = -0.131$). Thus, these two concepts—the proximity-bootstrapped equipment and the fail-open/safety wristband—are a reasonable choice to put together. Between the two of them, 7 out of 11 participants have at least one choice of

a system that they liked, 6 of the 11 participants have a system that they said that they would select, all 11 can have at least one system that they do not dislike, and 8 of the 11 have a choice of two systems that they do not dislike.

Satisfying the Stragglers: UV-Visible Tattoo. Despite its poor overall ratings (see Table 2.2), if we were to add a third system to the solution portfolio, the best choice would be the UV-visible tattoo. The criticality-aware IMD has slightly better ratings overall, but its ratings are highly correlated with the proximity-bootstrapped equipment. On the other hand, both of the participants who like the UV-visible tattoo liked neither the proximity-bootstrapped equipment nor the fail-open/safety wristband. The like/dislike ratings for the UV-visible tattoo have slightly negative correlations with the like/dislike ratings for both the proximity-bootstrapped equipment ($\tau = -0.196$) and the fail-open/safety wristband ($\tau = -0.125$), which indicates that it might be a useful addition to the solution portfolio. While the tattoos are not particularly popular overall, the UV-visible tattoo seems to pick up a segment of people who are not served by the other solution options. In fact, given a choice of these three systems, 9 out of 11 participants have at least one system that they like. The remaining two participants did not like any of the systems, so no combination of systems would include them.

A Portfolio of Systems

It is interesting to note that the above solution portfolio includes one system concept from each of the three different categories of security approaches. Like/dislike scores within each category tend to be somewhat highly correlated (for example, criticality-aware IMD vs. proximity-bootstrapped equipment, $\tau = 0.510$; fail-open and patient-specified functionality wristbands, $\tau = 1.000$; fail-open/safety wristband vs. each of the other two wristband systems, $\tau = 0.545$). Thus selecting multiple systems from the same category would generally be redundant. To obtain broad coverage, it is necessary to select a diversity of types of security approaches which will appeal to

different patients.

2.3.3 Results: System Concept Properties

Although the population size was too small to generate definitive statistical analyses, the interview responses do provide qualitative insight into the perspectives and concerns of patients with implanted cardiac devices. Below is a list of (non-mutually-exclusive) properties—both desirable and undesirable—that surfaced during discussions with participants about the system designs.

Security. Participants expressed liking systems because of their perceived security benefits and disliking systems because of their perceived security flaws. Moreover, sometimes the same system received both positive and negative feedback on its security properties. As an example, different participants complimented and critiqued the medical alert bracelet system based on its security properties. Some participants appreciated that access to their IMD would be protected by a password, while others objected to the fact that their passwords would be carried around with them and might be acquired by others.

Access. Participants were very interested in how the system concepts affected their safety in emergency situations. One of the better-liked systems, the fail-open/safety wristband (45%, 5/11), was much better liked than the other wristband versions. Comparing the scores on the like/dislike scale for the fail-open/safety wristband system versus the other wristband systems using a Wilcoxon signed-rank test yields a test statistic of $Z = -2.121$ and p -value of 0.034, indicating that like/dislike scores were significantly higher for the fail-open/safety wristband system. The fail-open/safety wristband differed from the others by offering two features that enhance the safety of the patient. The fact that this wristband version stood out from the others suggests that participants valued it for its safety features.

Participants rejected some systems based on their perceived safety flaws. In almost all of these cases, the participants were worried that hospitals might not have the

correct equipment, causing their IMDs to remain inaccessible in an emergency. These fears were expressed about scanning equipment for the tattoos, a black light for the UV-visible tattoo, and the proximity-bootstrapped device. Further fears were expressed that medical staff might not know to look for or might not be able to locate a UV-visible tattoo. Additionally, some participants were concerned that the criticality-aware IMD did not have a manual override.

Privacy. Some system concepts were called out as having negative effects on privacy. In particular, participants worried that wearing something or having some other visual indicator would force them to broadcast their medical condition to others. In the words of one patient:

Subject E: *I don't like the idea of wearing the wristband...I already have a defibrillator. Why do I have to wear something on my hand...to show that I have-, that I have a defibrillator, that there's something wrong with me. No.*

Aesthetics. Participants disliked some system mockups because they found them to be unaesthetic. One participant commented that the tattoo was visually too “busy.” There were also frequent comments about the appearances of the wristbands. The wristbands are meant to be worn at all times; some participants saw their unattractiveness as a major obstacle to system adherence.

Psychological welfare. Participants disliked systems that they deemed to be psychologically distressing or not respectful of their personal dignity. In particular, participants stated that wearing or seeing something that would remind them of their condition could be upsetting. One participant objected to the medical alert bracelet on these grounds:

Subject M: *It would make me feel like an invalid...That I had this thing, like the Scarlet Letter or [laughs].*

Another participant felt that the fail-open/safety wristband would be very distressing if its alarm suddenly went off in reaction to a magnet.

Convenience. A prominent objection to the wristband systems was their inconvenience. This included both mental and physical inconveniences: the mental inconvenience of remembering to charge the wristband and keeping track of it; and the physical inconvenience of having something on one's wrist that could catch on objects, needs to be taken off when showering, etc.

Cultural and historical associations. Many participants had objections to the tattoo systems.

Subject M: *Well, I mean for-, because I'm Jewish it-, I'm not-, a tattoo on the arm to me means a concentration camp. So right away that's the immediate horror.*

In one case, a participant disliked the system because she associated tattoos with drunks. Clearly, it is not desirable to have a security system for a beneficial medical device to have negative associations of these types in patients' minds.

Self-image and public persona. Interestingly, one patient did not object to the visible tattoo system in principle or due to any personal associations; instead, she objected that having a tattoo would present a persona to others that would be inconsistent with the one that she wished to project.

Autonomy and notification. Some participants had strong negative reactions to the criticality-aware IMD based on the fact that it silently changes its mode in an emergency to give all programmers access. While this objection could be addressed by adding in an audio or vibrational notification, it is noteworthy that these participants felt so strongly about not being informed. Similarly, some participants appreciated the proximity-bootstrapped device because—assuming that the patient is conscious—the system involves implicit consent.

2.3.4 Results: Attitudes Towards Wireless IMD Security

Participants were asked a series of questions about the importance of safety, privacy, and health. Responses for each question were coded on a 5-point scale of -2 (strongly disagree) to 2 (strongly agree). Some responses were uncodable. All 13 participants agreed or strongly agreed that they were concerned with maintaining their health, with a mean score of 1.38. Most participants also agreed that they were concerned about the safety and privacy of their electronic information (mean score 1.00, 10 out of 12 agreed or strongly agreed), their personal privacy (mean 0.77, 10 out of 13 agreed or strongly agreed), and their physical safety (mean 0.82, 9 out of 11 agreed or strongly agreed).

On the other hand, a majority of participants disagreed when asked if they were concerned that someone might change the settings on their IMD without their permission (mean score -0.92 , 10 out of 12 disagreed or strongly disagreed) or that medical staff would be unable to change the settings on their IMD in an emergency (mean -0.80 , 7 out of 10 disagreed or strongly disagreed). Despite this apparent lack of concern about the security of their IMD, participants tended to agree that something should be done to protect the security of future IMDs (mean 0.89, 7 out of 9 agree or strongly agree). To illustrate some of these points of view, below are quotes from two participants who lie at the opposite ends of the spectrum.

The following participant did not want random, unauthorized parties to be able to access his pacemaker. He indicated that leaving open access means that some malicious party will attempt to take advantage of that opportunity:

Subject K: *If, if anyone el-, everyone else can do it, they will do it...Or someone will do it.*

In contrast, this participant was unconcerned about the possibility of a cyber-attack targeting IMDs.

Subject D: *I'm not gonna-, I think it's ridiculous to worry about the security of it...Anybody that wants to get to me that bad, be my guest.*

2.3.5 Discussion

Although the results from the previous section suggest that patients may be served by providing a range of options, the landscape is complex and consists of more than patients' preferences. HCI research can contribute to the formation of a security ecosystem for IMDs, but patient values and preferences must be weighed against other important constraints. There are several reasons why a single-system solution might be preferable over multiple options.

Mental stress and complications of choice. As my collaborators and I observed in the interviews, different people prefer different levels of involvement in their medical decisions; some people delegate decisions to their doctors while others are very involved in the process. Having different choices for IMD security could potentially create friction between doctors and patients if a patient desires an option that a doctor believes to be unsuitable. Alternatively, if doctors offer patients a free choice of several systems, patients might experience stress while deciding and be doubtful about the wisdom of their final choice.

Medical Ecosystem. This study of patient perspectives and preferences—while useful—does not provide complete context for the design of security systems in the IMD space. The medical ecosystem is a complex space with tangible successes and losses; it is neither prudent nor realistic to treat it like a consumer electronics market. For example, while offering multiple security system options might facilitate patient preferences, this suggestion ignores a variety of potential complications, including: speed-of-use and simplicity in emergencies and other time-critical environments; the costs of regulatory approval; the burdens of training; and the space, time, and money costs of acquiring and maintaining any necessary equipment.

The provider study, which is discussed in the following section, provides a coun-

terpoint to the patient study in terms of both methodology and targeted stakeholder group. The findings from the provider study can also be used to consider some of the needs and constraints of the medical ecosystem touched upon above.

2.4 Medical Provider Perspective

A patient's medical care is an ongoing process that is affected by regulation, device manufacturers, federal testing, insurance companies, hospital equipment purchases, primary care staff, specialist nurses and doctors, emergency care staff, operating room staff, and others. While the patient study provides some insight into the values and priorities of patients who live with implantable cardiac devices embedded in their bodies (see Section 2.3, [23]), the question of how medical providers perceive these technical computer security directions has not been addressed. Yet, to be effective, security must work for and with all key stakeholders. In the case of implantable cardiac devices, this includes not only the patients, but also the medical providers who—in one way or another—ensure that the implantable devices function properly and improve patient health.

The following study with medical providers builds on the patient study by investigating similar security system designs concepts. However, as described in the methods, the participant pool and study format differ.

My collaborators and I conducted security-oriented Envisioning Workshops (see Section 2.4.1) with a variety of stakeholders involved in the care of patients with implantable cardiac devices including: nurses, emergency physicians, cardiologists, anesthesiologists, and device manufacturer representatives. We present results on: (1) what participants find important with respect to providing care and performing their jobs; (2) the metaphors participants use to describe implantable cardiac devices and security systems for these devices; (3) participants' evaluations of potential systems that represent different directions in technical security design; and (4) participants' opinions on what security system properties should be sought or avoided due to

Group	Male	Female	Total
I	4	6	10
II	4	3	7
III	5	2	7
Total	13	11	24

Table 2.3: Number of workshop participants by group and gender.

domain-relevant negative side effects. To be clear, the purpose of this research is not to gather participants’ feedback on the security performance of these systems—after all, the participants are not security experts—but rather to gather information about how different access control systems might impact participants’ jobs and their ability to care for patients.

2.4.1 Study Design

Participants

In this study, my collaborators and I sought to investigate in detail the values, priorities, constraints, and themes that emerge in a complex domain. We conducted three workshops with medical providers in the United States: one in a large city on the west coast (Group I) and two in a large city on the east coast (Groups II and III). Participants were recruited through a snowball method in which the research team first sent email letters to previous contacts in the medical provider community requesting suggestions for potential participants and relevant mailing lists; the researchers then followed up on those suggestions with email letters of invitation to participate in the research. Recruitment efforts were initially extremely slow; this was partially because my collaborators and I needed to obtain permission from appropriate authorities (i.e., “gatekeepers”) and partially because we needed domain insiders to explain the importance of—and cultivate enthusiasm for—study participation (i.e., “advocates”).

I applied for and obtained approval from the University of Washington’s human subjects review board. In order to synchronize study protocols across the multiple institutions involved in this study, it was necessary to submit multiple modifications. Participants were compensated \$200 for their time; while this amount may seem unusually high, it was deemed appropriate in the context of the particular participant pool.

Across the three groups, a total of 24 medical providers (age: average=39, min=28, max=64, mode=31) participated in the study. Table 2.3 provides a breakdown of participant gender by workshop. Participants reflected a broad spectrum of roles in the medical ecosystem including: cardiologists and electrophysiologists (n=2), nurses and nurse practitioners in cardiology and electrophysiology (n=5), anesthesiologist (n=1), emergency physician (n=1), other physicians (n=2), physician assistant (n=1), medical residents (n=4), medical device manufacturer representative (n=1), biomedical informatics researcher (n=1), and venture capitalist (n=1).

Workshop Format

My collaborators and I wished to elicit participant values, priorities, and constraints for the security of implantable cardiac devices. We sought a method that would provide opportunities for open-ended ideation about device security as well as focused reactions to potential early-stage security concepts. We drew inspiration from and adapted Kensing and Madsen’s techniques for “generating visions” [52]—which integrates metaphorical design with a Future Workshop (particularly the critique phase)—and from Yoo et al.’s Envisioning Workshop [97], which emphasizes surfacing value tensions between diverse stakeholders. In addition, my collaborators and I sought both to collect individual reflections and to benefit from group discussion; thus, data collection included individual written materials as well as verbal group interactions. The workshop protocol is described below.

Each session lasted a total of two hours. Audio recordings were made of each

session and then later transcribed for analysis.

Implantable Cardiac Device Overview and Initial Perspectives. To ensure that all participants had some shared vocabulary for implantable cardiac devices, a research team member provided a brief overview of implantable cardiac devices and clarified how terms would be used during the workshop. This overview did not include information on security for implantable cardiac devices. Following this overview and to tap into participants' perspectives prior to any influence from the workshop activities, participants were asked to complete a brief paper and pencil worksheet that elicited their initial views on security and access control for implantable cardiac devices. The worksheet contained the following questions:

1. *What properties about implantable cardiac devices or the ecosystem surrounding them do you value most?*
2. *What things about implantable cardiac devices or the ecosystem surrounding them should not change?*
3. *What things about implantable cardiac devices or the ecosystem surrounding them most need improvement?*
4. *What is the most common problem related to implantable cardiac devices that you encounter in your line of work (e.g., lack of access to patient information, inability to access cardiac device, device malfunction)?*
5. *What is the problem with the most negative health impact (related to implantable cardiac devices) that you encounter in your line of work?*

Metaphor Generation. To help understand the broad backdrop of participants' perspectives as well as potential mental models, participants were invited as a group to share verbally: (1) metaphors for implantable cardiac devices; and (2) metaphors



Figure 2.2: Photos of system concepts used in the medical provider workshops.

for security and access control for those devices. A research team member facilitated the contributions and recorded each metaphor in a few concise words on a whiteboard.

Critiques and Concerns. To understand how security and access control systems for implantable cardiac devices could go awry as well as to understand medical providers' hesitations and concerns about this type of technology, participants were invited as a group to share verbally their concerns and fears about security for implantable cardiac devices. Volunteer participants grouped the concerns into clusters based on similarity.

Evaluation of Security and Access Control System Concepts. To understand participants' views on what properties to advocate for and which to avoid in the development of security and access control solutions for implantable cardiac devices, a researcher introduced participants to six potential security and access control

systems, one at a time. The researcher indicated that these were early, representative systems designed to elicit feedback. The system concepts are described in Section 2.2. For this study, my collaborators and I chose to include the centralized database system concept, exclude the visible tattoo of a password, and present only the fail-open/safety variant of the wristband (i.e., Systems I-II, System IV, System V' from Section 2.3.1, Systems VI-VII). Figure 2.2 shows the photos of the system concepts that were shown to providers during the workshop. For each system, participants completed a paper and pencil worksheet in which they recorded their responses to the following questions: *From your perspective as a professional who deals with implantable cardiac devices, what do you like about this concept? What do you dislike about this concept? Why?*

Once participants had been introduced to all six system concepts, participants completed a worksheet with the following questions:

1. *Would you say that you like any of the concepts, and if so, which ones?*
2. *Would you say that you dislike any of the concepts, and if so, which ones?*
3. *If you were to choose one or more of these concepts to recommend for use in the future, which concept or concepts would you choose? Why?*
4. *If you were to choose one or more of these concepts to recommend against use in the future, which concept or concepts would you choose? Why?*

Open-ended Discussion. Finally, to ensure that participants had ample opportunity to surface any major issues that might have been missed, participants engaged in an open-ended discussion around security and access control for implantable cardiac devices in which they could respond to and debate each other's ideas. To initiate the discussion, the workshop facilitator used the following prompt: *What are the challenges in this space?*

Coding and Reliability

Participants' written initial perspectives were coded systematically using the following process. One researcher developed a coding scheme using all of the data; once completed, that researcher used the finalized coding scheme to systematically recode the entire data set. A second coder—not affiliated with the research team or the study—was trained in the coding scheme using data from 4 participants, and then independently performed reliability coding of the data for the remaining 20 participants. This process resulted in an overall kappa of 0.75; Fleiss rates any value of kappa over 0.75 as excellent agreement and between 0.40 and 0.75 as intermediate to good agreement [31], while Landis and Koch rate a kappa of 0.81 to 1.00 as “almost perfect” and between 0.61 and 0.80 as “substantial” agreement [55].

The metaphor data set was smaller and, thus, more appropriately coded by consensus. My collaborators and I used the following process (1) first, two researchers independently read through all of the data to generate an initial set of coding categories and assign responses to categories; (2) next using consensus, researchers iteratively synthesized categories and arrived at agreement; and (3) then both researchers made another independent pass through all of the data and any lingering disagreements were resolved.

Justifications in the security system concept evaluation data were identified from inspection of the qualitative data and are presented via participant quotes.

2.4.2 Results

Given the relatively small number of participants in each workshop, there was no way to draw meaningful comparisons among the workshops' participant demographics (e.g., location, gender, age, professional role). The data from all three workshops was combined into one data set.

Initial Perspectives

Participants' written responses to the Initial Perspectives questions provide a relatively unbiased (that is, largely uninfluenced by our subsequent workshop activities) view into what participants consider important about implantable cardiac devices and their usage to treat patients. Since our primary interest was to understand broadly the pre-existing issues important to medical providers, I examined providers' responses to the set of five questions as a whole (rather than by individual question).

Thirteen categories of issues emerged from the analysis of participant responses as follows (in alphabetic order): (1) **Access & Sharing**; (2) **Compatibility**; (3) **Correct Usage**; (4) **Device Battery Life**; (5) **Device Compactness / Inertness**; (6) **Device Ecosystem**; (7) **Device Functionality**; (8) **Patient / Patient Health**; (9) **Programming**; (10) **Quality of Data**; (11) **Remote Monitoring**; (12) **Security & Privacy**; and (13) **Surgery & Healing**. Table A.1 in Appendix A.1 provides definitions for the categories, example participant responses, and the percentage of participants who raised each issue.

Over three-quarters of the participants expressed issues related to Device Functionality (79%) and Patient/Patient Health (75%); and more than half mentioned Surgery & Healing (58%). The next most represented categories were mentioned by roughly a quarter of the participants (ranging from 25–29%). That said, given the sample size and exploratory nature of this study my collaborators and I believe it would be prudent to consider all 13 categories of issues when designing a security system for implantable cardiac devices. This list of issues provides a window into the values and priorities of medical-provider stakeholders in the medical ecosystem. Security and human-computer interaction researchers may not have sufficient domain knowledge to make direct judgments as to how a system design might interact with these aspects of medical care; these categories, and other data like them, may serve as a meaningful starting point for dialog with domain experts.

Metaphors

Metaphors often underlay people’s mental models of technological systems, which can affect the ways in which they interact with those systems (e.g., [34, 90]). The metaphors supplied by participants provide some indication as to how they conceptualize implantable cardiac devices and security systems for those devices. In addition, using metaphor generation as an opening activity was intended to help break the ice; metaphor generation is rapid and appropriate for ideas that might otherwise be considered offbeat or silly.

As a group, participants generated a total of 81 metaphors: 42 for the implantable cardiac devices and 39 for security and access control for those devices. The following 11 categories—given in alphabetic order—emerged from clustering together similar metaphors: (1) **Agency**; (2) **Bio-medical**; (3) **Business**; (4) **Emotion**; (5) **Information**; (6) **Maintenance**; (7) **Personal Identity**; (8) **Privacy**; (9) **Risk**; (10) **Security**; and (11) **Technology**. Table A.2 in Appendix A.2 provides definitions for the categories, example metaphors for each category, and the number of metaphors identified in each category.

Participants understood implantable cardiac devices in a broad variety of ways, including bio-medical terms, both positive (e.g., “life savers”) and negative (“site of infection”); and emotional terms, though always negatively (e.g., “anxiety producing,” “source of hassle”). Within these framings for the device itself, participants described the device’s security systems in terms of security, both secure (e.g., “secure site on the Internet”) and insecure (e.g., “bank with an unlocked vault”); risk, both mitigating (e.g., “insurance policy”) and vulnerable (e.g., “life threatening”); and information, both known (e.g., “complete control of information”) and unknown (e.g., “black box on a plane”). The diversity of metaphors as well the potential for any given metaphor to convey both positive and negative dimensions points to the need for security researchers to attend carefully to how stakeholders “conceptualize”

Providers N = 24	Participant Percentage			
	Like	Dislike	Recommend	Rec. Against
I. Medical alert bracelet with password	29	46	21	33
II. Centralized database	38	21	25	25
IV. UV-visible tattoo of a password	17	54	13	50
V'. Fail-Open/Safety wristband	58	17	46	13
VI. Proximity-bootstrapped equipment	38	25	38	21
VII. Criticality-aware IMD	38	42	33	38

Table 2.4: Percentage of participants by security system concept who liked, disliked, recommended or recommended against each system concept. Light shading indicates reasonably high satisfaction with a system concept; heavy shading indicates fairly low satisfaction.

in lay terms security for such devices, and how they use those conceptualizations to generate positive or negative perspectives on the security system.

Security System Concepts

To understand participants' views and values about the exemplar security system concepts, my collaborators and I first looked systematically at which systems participants found strongly acceptable—that is, the systems which many participants liked and very few participants disliked—and which systems participants found less acceptable—that is, the systems which few participants liked and many participants disliked. This data helps inform the interpretation of the following section, in which I present some of the reasons for which providers expressed liking or disliking the systems.

Evaluations. As with the patient study, these results give some indication of participants' relative preferences for the presented system concepts; however, these results should not be overemphasized. The methodology and sample size of the study are geared towards an inductive investigation of the design space and participants' reactions to system concept properties, rather than a deductive ranking of system concepts.

Table 2.4 provides an overview of the results from participants' evaluations of the early-stage security system concepts presented during the workshop. The fail-open/safety wristband (indicated in Table 2.4 with lightly shaded cells) was the best received in all categories: the largest percentage of providers liked it (58%) and would recommend its use (46%), and the smallest percentage of providers disliked it (17%) and would recommend against its use (13%). The UV-visible tattoo of a password (indicated in Table 2.4 with a row of darker shaded cells) was the least satisfactory in all categories, with only 17% of providers liking it, 54% disliking it, 50% recommending against its use, and only 13% recommending its use. Two other systems reach relatively high thresholds on dislike and recommend against: the medical alert bracelet with a password (46% dislike and 33% recommend against) and the criticality-aware fail-open IMD (42% dislike and 38% recommend against); these data might suggest avoiding the use of variants of any of those three system concepts.

In general, when one examines system evaluation results, one requires high satisfaction thresholds (i.e., high “like” and “recommend” percentages, low “dislike” and “recommend against” percentages) in order to describe a system as well-liked. In contrast, less stringent thresholds are necessary to describe a system as problematic, in order to respect the perspectives of stakeholders who may be in the minority. This is in line with the value sensitive design dams and flows analysis guidelines (e.g., [63]).

Justifications. As noted above, the evaluation results on the exemplar security system concepts immediately raise the question of why providers like or dislike a system or would recommend for or against its use. Recall that Table 2.1 provides a breakdown of some of the properties embodied by the various system concepts, such as requiring physical proximity to the patient or having a manual override. Providers are potentially reacting to these properties in their evaluations. Below I report what system properties providers said they liked and disliked about each system; for the systems that were particularly high-ranked or low-ranked, I break out the relevant properties as lists.

System I (Medical Alert Bracelet with Password). The medical alert bracelet was one of the system concepts most disliked (46%) and recommended against (33%). Providers most frequently expressed disliking System I for the following reasons:

- ↓ Access is not guaranteed—the bracelet may be forgotten, lost, stolen, damaged, or the patient may choose not to wear it. (E.g., “*In an accident, the bracelet could be damaged/lost and emergency personnel would not be able to access device*”)
- ↓ The security is insufficient. (E.g., “*Way too easy to maliciously steal password*”)
- ↓ It visibly indicates to patient and others that the patient has a condition. (E.g., “*Identifies pt. as having a problem*”)

The relatively poor reception of System I suggests that either these properties are particularly disagreeable to participants, or that the advantages are not sufficient incentive to tolerate the disadvantages. When participants expressed liking that the medical bracelet solution they noted that the system did not depend on other equipment or systems, provided “reassurance” to the patient, was cheap, and provided some security.

System II (Centralized Database). The centralized database was neither one of the highest-rated nor one of the lowest-rated systems. Participants expressed concerns about: the availability of the database across regions, across providers and manufacturers, in case of disaster, or in case of other technical difficulties; how to identify patients to look them up in the database; how to secure the database and identify who is authorized to access it; who will administer the database and how they will fund and maintain it; and the fact that a database would require time away from the bedside to access. In contrast, participants appreciated that this system

neither required nor depended upon the patient to wear anything, was theoretically universal, and provided more security than System I.

System IV (UV-Visible Tattoo of a Password). The UV-Visible Tattoo was the lowest-ranked system for all evaluation questions. Participants expressed disliking this system for the following reasons:

- ↓ Required equipment may not be working, accessible, or timely to acquire. (E.g., “requires UV light (*i.e.* working bulb, power source)”))
- ↓ Patients may have cultural, social, or personal objections over a tattoo. (E.g., “religious restrictions against tattoos”)
- ↓ Access is not guaranteed the tattoo could be faded, damaged, or distorted. (E.g., “blood or trauma may obscure tattoo”)
- ↓ Password revocation or changes could be complicated. (E.g., “how to change when device is changed out”)

Again, this suggests that the disadvantages outweigh the properties about the system which participants liked: its invisibility in the patient’s daily life, both for human and security reasons; and the fact that it is (theoretically) always with the patient, but requires no patient effort.

System V’ (Fail-Open/Safety Wristband). The fail-open/safety wristband was the highest-rated system across all categories. Participants reported liking the system for the following reasons:

- ↑ The fail-safe mode guarantees access. (E.g., “*GREAT* failsafe mode (remove bracelet)”))
- ↑ The system provides some safety features. (E.g., “safety features *BIG plus*”)

- ↑ The system provides some security features. (E.g., “*Provides mechanism against snoopsequivalent to locking your door when you leave the house*”)
- ↑ The mechanism gives access control power to the patient. (E.g., “*pt. feels empowered. pt. is an active participant in their own care*”)
- ↑ Provides a visual cue to EMTs. (E.g., “*identifies patient as having an ICD*”)

Following previous lines of reasoning, these advantages must outweigh the dislikes expressed by participants: that there is no security if the wristband is not worn (and it is easily removed); that the wristband requires battery replacements or recharging and requires the patient to wear it, for which there is no incentive; that there may be many false-positive calls to 911; that the patient is visibly identified as having a medical condition; that emergency medical staff would require training to know to remove the wristband; and that the system is potentially expensive to develop and produce.

System VI (Proximity-Based Equipment). The proximity-based equipment was neither one of the highest-rated nor one of the lowest-rated systems; this system most closely resembles the status quo of access control for implantable cardiac devices. Participants expressed liking a variety of system properties: that it provides some security from wireless tampering; that it does not require the patient to wear or do anything (and therefore does not provide a visual indication of the patient’s condition); that it does not depend upon other equipment or systems; that it is similar to the current model; that it is easy, and allows bedside access; that it would be a (theoretically) universal access system; and that it gives the patient some control over who may access their device. Conversely, participants reported disliking: that patients are still susceptible to in-person security breaches; that such a system would require new equipment, which is expensive; that such a system would potentially be manufacturer-specific; and that such a system would require all medical centers to

have the equipment on-hand and readily accessible for emergencies.

System VII (Criticality-Aware Fail-Open IMD). The criticality-aware IMD was one of the systems most disliked (42%) and recommended against (38%). Providers most frequently expressed disliking System F for the following reasons:

- ↓ The IMD may not correctly identify a medical emergency (false negative—closed access). (E.g., *“this assumes the device can properly recognize emergencies→current devices can’t even recognize some arrhythmias correctly”*)
- ↓ The IMD may incorrectly identify a medical emergency (false positive—open access). (E.g., *“possibility of misidentifying a medical emergency”*)
- ↓ There may be non-emergency situations where the IMD needs to be accessed. (E.g., *“what happens if the patient moves and has a new cardiologist?”*)
- ↓ This system could change IMD size or shape, consume battery life, or cost more. (E.g., *“will certainly add expense to cost of device such that CMS may veto payment/reimbursement”*)

As previously reasoned, these disliked properties apparently overpower the properties that participants liked: that it (theoretically) allows access in an emergency; that it provides some security; that it depends upon no extra equipment; and that it does not require the patient to do or wear anything.

2.4.3 Discussion

Recall that the medical provider study was conducted in part to complement the prior study on patients view and values about early-stage security solutions for implantable cardiac devices [23]. The set of systems presented to medical providers was the same as those presented to patients with the following exceptions: three systems

that were deemed no longer viable security concepts were removed (i.e., visible tattoo of a password, fail-open wristband, and fail-open/patient-specified functionality wristband) and one new security concept was included: a centralized database. In terms of format, the manner of presenting information about the security concepts and obtaining responses from tailored to the participants. For highly literate, professional medical providers security systems were presented in a group setting via a verbal description and supporting slides; and medical providers provided individual written feedback in a focus group setting. For patients—many of whom were elderly or ill—security systems were presented to each individual with a verbal description and a physical mockup as a prop; patients provided verbal feedback as part of their semi-structured (individual) interviews. The questions asked while similar in substance were also slightly different as appropriate to the individuals' role as a medical provider or patient. Specifically, while both providers and patients were asked if they liked/disliked any of the systems, and if so, which ones; medical providers were asked if they would recommend for or recommend against using any of the systems, while patients were asked if they would choose to use any of the systems.

I turn now to consider some of the results from both studies together as a way to explore security concepts that might be successful for both sets of stakeholder groups. Considering first the proximity-based equipment approach, in the previous patient study (see Section 2.3, Table 2.2), this security concept was least disliked (0% dislike), and hence might be the most logical system to choose; however, 25% of the medical providers disliked the proximity-based equipment approach and 17% would recommend against its use, making it a less desirable choice overall. In a similar vein, the criticality-aware fail-open IMD was more liked (27%) than disliked (18%) by patients; however it was more disliked (42%) than liked (38%) by providers. These findings suggest a different level of understanding, awareness, or requirements between patients and providers. In the case of criticality-aware fail-open my collaborators and I suspect that this difference is primarily due to providers' higher concern regarding

the lack of a manual override if the system fails to recognize a medical emergency. The provider results suggest that a criticality-aware fail-open approach may not be a suitable solution for securing IMDs.

In terms of similarity of perspective, the fail-open/safety wristband approach was the security concept that was least disliked by the medical providers (17% dislike, 13% recommend against), it was also the most liked by the patients (45% like), and also most liked (58%) by medical providers, with 46% recommending its use. The UV-visible tattoo system was more disliked than liked by both groups (27% dislike and 18% like among patients; 54% dislike and 17% like among providers).

The patient study recommended a set of three solution choices that, if offered together, might satisfy the desires of most patients: a proximity-based equipment system, a fail-open/safety wristband, and a UV-visible tattoo of a password. Given the strong opposition to UV-visible tattoos among providers (50% would recommend against), however, the workshop results caution against their use in practice.

2.5 Recommendations for Design

As previously indicated, while the patient and provider results do contain quantifiable evaluations of the security system concepts, the primary purpose behind these more inductive study designs is to obtain more qualitative information about the needs and restrictions of stakeholders in the target domain. The qualitative results from the studies help produce a set of concrete recommendations and priorities for researchers to consider when designing security systems for implantable cardiac devices.

2.5.1 Patient-Centric Recommendations

Based on participants' reactions to these specific systems as instantiations of general system properties, if researchers and designers want their IMD security systems to be liked by—or at least acceptable to—patients, my collaborators and I advise them to meet the following criteria:

- **Good (perceived and actual) security properties.** Some participants objected to the medical alert bracelet system because they felt that carrying around a human-readable password was insecure.
- **Good (perceived and actual) access properties.** Participants were very interested in whether a system offered reasonable access: for example, whether it had an override or whether medical staff might not have necessary equipment.
- **Respect patients' privacy and avoid disclosing patients' conditions.** Both the medical alert bracelet and the wristband systems were criticized for being visible indicators of patients' medical conditions, while the emergency and warning wristband was criticized because its alarm feature would require explanation if it went off in a crowd.
- **Be aesthetically pleasing (or at least aesthetically neutral).** Many participants objected to the wristbands based on their appearance.
- **Avoid causing patients sudden alarm about their health.** A participant expressed that an alarm suddenly going off might cause patients to panic.
- **Avoid needlessly reminding patients of their condition.** Depending upon the medical condition and the implanted device, patients may need to maintain some awareness of how their behavior might affect their health; nevertheless, the psychological effects of the technology should be minimized. Some participants did not want to look at medical alert bracelets or wristbands that would remind them of their conditions.
- **Avoid being physically irritating.** The medical alert bracelet and the wristbands were both criticized because they could be physically irritating to wear on the wrist.

- **Avoid requiring frequent upkeep.** The wristbands were also criticized because they require regular recharging.
- **Work with patients to offer an option that fits their self-image.** One participant in particular expressed that the visible tattoo was not palatable because it would give other people a certain impression about her—an impression that she did not want to project. Systems should be mindful of patients' sense of dignity.
- **Avoid unwanted negative associations due to historical, religious, or cultural factors.** Several participants strongly disliked one or both of the tattoo systems because of tattoos' associations with concentration camps; another participant did not like the visible tattoo system because it reminded her of drunks.
- **Provide the patient with the option to be notified of changes in system status.** Several participants disliked the criticality-aware IMD because it changed its access mode without notifying the patient.

2.5.2 *Provider-Centric Recommendations*

Drawing on a synthesis of the results and insights from immersion in the study data, I now provide concrete design considerations for security researchers working in the implantable cardiac device domain.

Access, Access, Access

Access—and the related issue of compatibility—show up in both the initial perspectives and the evaluation of security system concept data sets, and are particularly stressed in the latter. Participants repeatedly indicated the importance of access along a variety of axes:

- Providers must always be able to access the implanted device, and security systems should either fail to an open state or offer some kind of override.
- “Unplanned” access does not only occur in emergencies; patients may travel or change cardiologists, and records are not always transferred smoothly.
- Access should not rely upon a centralized system, which could be unavailable (due to technical, geographic, or other reasons) and which merely defers the security problem.
- Access cannot rely upon a conscious or compliant patient.
- Access should avoid relying on additional equipment, which can delay or block patient care or remove providers from the bedside.
- Access should be timely, and should therefore require few steps. Perhaps, above all, in the words of one of the participants: *“Please, please, please keep it SIMPLE.”*

Mechanics and Logistics

Various aspects of IMD mechanics and logistics are raised in both the initial perspectives and security system concept evaluation data sets, and any security system should avoid disturbing the status quo in terms of:

- cost, which can also affect insurance approval;
- required training, particularly for non-cardiology staff;
- implant battery usage;
- implant size; or

- any other aspect that might impact the surgical or healing processes.

Safety Features and Incentives

Participants showed interest in the possibility of incorporating safety features into a security system for IMDs. The exact nature of such features and how they might be tuned should be further investigated; for example, many participants expressed concern that a “911 feature” would result in many false-positive emergency notifications. Advantageously, depending upon the system design such safety features might provide incentives for patients to use the system. This is particularly relevant in the case of a system like the wristband (System V'), with which the patient only receives security if they choose to wear the band.

Empowering Without Burdening

Ideally, patients should be given some implicit or explicit role in the access control process, whether via overt action or by allowing someone extended skin contact. Generally speaking, such a role might give patients a feeling of empowerment, but more practically speaking, patients could provide human reasoning as to whether or not their device should be accessed in a given situation. Conversely, patients should not be unduly mentally or physically burdened by a security system. As one example of this, anything that visually indicates the patient's condition should be opt-in; visual indicators such as medical alert bracelets are useful to emergency staff, but patients should be able to weigh the advantages and disadvantages and choose whether or not to participate. Moreover, this consideration raises a slew of ethical, legal, and philosophical questions: Should a security design hinge upon patients being able to choose whether or not they wish to comply? How many patients would actually comply? Should a security design strive to equally protect all patients from potential harm, or is that attitude paternalistic? What are the repercussions, legally or in

terms of reputation, if a company's IMD is attacked, and security was optional? The domain is full of interesting questions that are ripe for investigation.

2.6 Summary

The work reported in this chapter makes three important contributions. First, I have provided detailed case studies that demonstrate how security researchers can draw upon diverse direct and indirect stakeholders to understand the relevant properties of a technology domain—particularly one that is either not well established or not well understood outside its field. This work points to value tensions both within a stakeholder group—as evidenced by disagreement among medical providers about which security concepts are preferred—and between stakeholder groups, as evidenced by disagreements between the medical devices providers and patients—again, regarding security system concept preferences. While these cases were focused on implantable medical devices, the methodologies could be applied to a range of other emerging technologies.

A second contribution concerns method. Specifically, I foregrounded early-stage security systems in order to gather explicit feedback on potential security directions and to identify value tensions. Part of the methods' effectiveness comes from developing clear, meaningful ways to convey complex security concepts to interview and workshop participants through the use of concrete system explanations that embody the essence of the security solution. Outside of this common theme, the two studies utilized different methodologies in order to adapt to the different characteristics of the participant population. In the case of the cardiac patients, my collaborators and I developed a semi-structured interview that allowed us to adjust pacing and explanations. For the medical providers, we adapted an established method used in value sensitive design, the Envisioning Workshop; this format allowed us to capitalize on the presence of people in multiple professional roles and initiate group dialogue and debate.

As a third contribution, this study offer domain specific findings for implantable cardiac device security—a topic of interest within the security community. Security experts can utilize the data from this study to inform the design of security systems, with the goals of increasing system adoption, supporting correct usage of security systems, and avoiding negative system side effects.

Bridging between early-stage technical innovation and the lives of stakeholders who will be impacted by that technology downstream is not easy; however, these connections are essential to enable technologists to do work that is informed by the values, priorities, and constraints of the people for whom they are ultimately designing. Toward that end, the work reported here provides case studies to suggest how such work could be done, and methods for making progress towards incorporating human values into technical security design.

2.7 Acknowledgments

This work was supported in part by the US National Science Foundation (NSF) under awards CNS-0846065, CNS-0905118, and CNS-0905384, by the US Department of Health & Human Services (HHS) under award HHS-2010-03958-09, and by an Alfred P. Sloan Research Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this dissertation do not necessarily reflect the views of the funding agencies.

I thank the patients and medical providers who participated in my study. I thank the clinic staff who helped facilitate interview logistics. I thank Miro Enev, Daniel Halperin, Karl Koscher, Benjamin Ransford, and Alex Takakuwa for contributing to the data analysis process.

Chapter 3

AUGMENTED REALITY AND BYSTANDER PRIVACY: EXPLORING PRIVACY ATTITUDES WITH AN EMERGING TECHNOLOGY

The work in this chapter of my dissertation deals with understanding people's privacy perspectives on an emerging technology that, unlike implantable cardiac devices, is only beginning to come to market: augmented reality glasses. These devices are interesting to study because they are wearable, incorporate audiovisual recording, and are not currently prevalent. In particular, my work with augmented reality devices has focused on how bystanders to these technologies perceive the impacts to their privacy. The results from this work, as with my work with implantable cardiac devices, help designers and technologists understand people's objections to a technology and support design decisions that address people's concerns. In this chapter, however, the investigations are centered around concerns regarding how a technology might violate people's privacy, rather than around how a technology's security system might negatively impact surrounding logistics and values. Additionally, the methodology used in this chapter (in-situ interviews) differs from those used in Chapter 2 (semi-structured interviews with mockup props and group workshops); in-situ interviews were used to help ground study results for a technology that is not currently prominent, and with which participants might have limited experience.

Section 3.1 provides some background on augmented reality technologies and research dealing with audiovisual recording and privacy issues. Section 3.2 provides the study design, results, and discussion of in-situ interviews with participants regarding their perspectives on augmented reality devices and recording. Section 3.3 lays out

design axes for privacy-mediating approaches that could be used in this space; providing a more structured overview of the design space helps guide design decisions and reflect on the current body of research. Section 3.4 summarizes this chapter in the context of my dissertation.

Part of the material in this chapter first appeared in [24]. Collaborators on this work included Zakariya Dehlawi and Tadayoshi Kohno.

3.1 Motivation and Overview

Audiovisual recording is pervasive in public spaces. This recording takes place predominantly via two classes of devices: handheld devices such as camera phones, and infrastructure devices such as closed-circuit television (CCTV). These two recording paradigms can be characterized and contrasted via axes such as mobility, recording cues, typical recording duration, content ownership, and intended usage.

A new form factor for recording hardware—glasses-style augmented reality (AR) devices—is poised to become more common. If commercialization attempts succeed in creating a market for these types of devices, there could be a massive increase in the number of people using wearable cameras. This class of device shares characteristics with both camera phones and CCTVs; however, the result is a unique amalgamation of properties. For example, AR-style glasses—unlike camera phones—are well-suited for periodic, continuous, and low-effort audiovisual recording. In contrast to CCTVs, AR glasses are mobile and controlled by individuals.

While research has been conducted on the relationship between recording and privacy, most prior work focuses on the current dominant form factors. There is a need for more research into how wearable and glasses-style devices differ from other classes of cameras. Moreover, these cameras have not yet achieved significant market penetration. As a result, we have the opportunity to study how perceptions and usage patterns change over the adoption of a new technology.

In this study we consider the perspectives of bystanders of AR glasses. In partic-

ular, I consider their perceptions of how and why these devices might impact their privacy. Bystanders are particularly relevant for study, as they are the largest stakeholder group—even users themselves are bystanders to other AR devices. We have the opportunity to explore technology designs that can mitigate bystander concerns.

In this study my collaborator and I report on our in-situ approach to investigating bystander perspectives on AR-style recording. We wore a mock AR device in cafés around a city over the course of 12 field sessions. During these sessions, we conducted semi-structured interviews with 31 individuals on their reactions to the co-located device. The contributions of this work are as follows:

In-Situ AR Perspectives From Bystanders. This work provides the first (to my knowledge) in-situ look at people’s perceptions of, and reactions to, glasses-style AR devices. This analysis of interview data surfaces: (a) reasons why participants do or do not consider AR glasses to change the bystander experience; and (b) factors that contribute to participants not wanting to be recorded. Additionally, I explore participant thoughts on permission and blocking technologies for recording.

Design Axes for Privacy-Mediating Technologies. In parallel, I use an exploration of the background literature to formulate design axes for privacy-mediating recording technologies.

3.1.1 *Related Work*

Infrastructure Recording Technologies

Early research on media spaces—such as by Adams [6] or Bellotti and Sellen [12]—explores the privacy issues that result from an environment instrumented with recording capabilities. This research is particularly transferable to CCTVs, but also has some transferability to AR privacy issues for bystanders.

More recently, Nguyen et al. interviewed participants to explore their feelings about CCTV recording [65]. They interpreted their results largely via Smith’s Con-

cern for Information Privacy model [82]; this model breaks privacy concerns into the dimensions of collection, improper access, unauthorized secondary use, and errors.

Massimi et al. use the Day Reconstruction Method (e.g., [51]) to interview participants about the recording technologies that they encounter in their daily lives [60]. Their results have a heavy focus on infrastructure-style CCTV cameras rather than on individuals' mobile cameras.

Friedman et al. conduct an in-depth analysis of interviews with bystanders to a camera installation recording a public fountain area [36]. The authors investigate underlying issues and interviewee justifications. For example, participants viewed the installation to be less acceptable if the footage was streamed to a remote location.

Mobile Recording Technologies

Steve Mann (e.g., [9]) and Thad Starner (e.g., [85]) have bodies of work on AR technologies. More topically, they have both worn AR devices for extended periods of time and in public. They have anecdotally reported their experiences wearing AR devices. For example, in 2012 Mann reported that he was assaulted by a staff member in a Paris McDonald's due to his use of EyeTap [58].

Nguyen et al. conducted a study with many parallels to the work in this chapter [66]. They also wished to investigate bystander reactions to a wearable camera. However, the camera in question was one primarily positioned as an assistive device for users with memory or vision impairment. The stated purpose potentially affects bystander reactions to the device. The study collected data via paratyping (see below).

Methodologies

Paratyping (e.g., [5, 49]) is a methodology for collecting in-situ feedback from bystanders via situated experience prototyping. With this technique, participants are recruited to act as proxies for the researchers. Participants carry short surveys and



Figure 3.1: A photo of the mockup device worn during field sessions. The device is a Myvu Crystal Personal Media Viewer with an attached (non-functioning) camera.

distribute them to bystanders with whom they interact. The surveys, which are returned by mail, probe bystander reactions to ubiquitous technologies in the context of their recent interactions. These bystanders can optionally be contacted for follow-up interviews.

Choe et al. investigate participant attitudes to sensors in the home via sensor proxies [19]: participants placed repackaged motion sensor lights in locations around the home. Light activation served to probe participants to record reactions to the hypothetical sensors in study diaries.

Mancini et al. explore reactions to hypothetical technologies via a video prototyping methodology they call ContraVision [57]. In particular, they present a video that portrays the technology in a positive light together with a video that portrays the technology in a negative light.

3.2 In-Situ Interviews with Bystanders

3.2.1 Study Design

Field Sessions and Interview Protocol

Each field session was conducted using two researchers: a researcher wearing a mock AR device and a researcher conducting interviews. The AR mockup consisted of a pair of media glasses—the Myvu Crystal Personal Media Viewer—and an attached, non-functioning camera (see Figure 3.1); see Section 3.2.1 for the reasoning behind the decision not to record. Field sessions proceeded as follows: the interviewing researcher would enter the café, order food or drinks, and take a seat. The researcher wearing the mockup device would then enter the café, order, and sit down to work. Patrons that had obvious reactions to the AR device, were likely to have noticed the AR device, or were pertinent for theoretical sampling (e.g., were accompanied by children) were approached for an interview. The interviews were semi-structured and based around the following questions:

1. *Did you notice the glasses that (s)he's wearing? What about them did you notice?*
2. *Have you heard about those kinds of glasses? What have you heard?*
3. *Did you know that those kinds of glasses have electronics and a display attached?*
4. *Did you know that you can record video with those kinds of glasses?*
5. *Why do you think someone would want to wear those kinds of glasses?*
6. *Do you think recording with those glasses is similar or different to recording with a cell phone? Why?*

7. *How do you feel about being around someone who is wearing those kinds of glasses? Why?*
8. *Would you want someone with those kinds of glasses to ask your permission before recording a video?*
9. *Would you be willing to wear something that would block someone from being able to record you?*

The progression of increasing specificity in the questions was arranged to probe participants on the more general topics in a non-leading way. The protocol served as a guide for the interview; questions were modified or discarded based on the flow of the conversation and any time constraints set by the participants. Interviews were not recorded: the interviewing researcher took interview notes and both researchers made observation notes. The researchers—one male and one female—took turns wearing the AR mockup and conducting interviews. The human subjects Institutional Review Board of the University of Washington reviewed and approved the study protocol.

Methodology Discussion

The investigative methodology has both benefits and drawbacks. My collaborators and I deploy (non-recording) AR-style glasses into real environments and give participants a chance to observe them before they are interviewed. During interviews, participant responses are grounded by the presence of the device in the environment. Moreover, because we approach individuals “in the wild,” we are potentially able to interview people who do not respond to research recruitment ads.

My collaborators and I found cafés to be suitable settings for a number of reasons: they are publicly accessible; they have a reasonable throughput of traffic; they are settings where researchers can position themselves for extended periods of time; and they are environments in which it is plausible to approach individuals for an interview.

Moreover, within a single city—and even within a neighborhood—the character of a café and its clientele can vary greatly. Different cafés attract different demographics and subcultures. At various times, cafés draw people engaged in a variety of activities: socializing, eating, reading, meeting, working alone, studying in groups, or playing games.

As mentioned, this investigative methodology has some drawbacks. My collaborators and I chose not to record interviews. This was done both to make the process less daunting to participants and to facilitate soliciting perspectives from individuals who might object to the idea of being recorded. Although the interviewer took notes during and after the interview, the pace of the interview and the need to engage with participants inevitably means that these notes are not as complete as a full transcript. Additionally, my collaborators and I attempt to be respectful of potential participants and their time. The interviewer chose to approach only those individuals who seemed like they could be interrupted; this meant, for example, that individuals focused on their laptops and groups in deep conversation were excluded.

Data Collection

A total of 12 field sessions were held in 8 different cafés over the course of 3 1/2 months in spring and summer 2013, and ranged in duration from 20–90 minutes each. The field sessions were performed at different times of day and on different days of the week, including weekends. At the end of an observation session, individuals or groups were solicited for interviews. If a group was approached, everyone in the group was included in the interview. These 12 field sessions yielded 23 interview sessions with 31 participants. The participants (M=18; F=13) represented a variety of age groups (18–22=8; 23–25=5; 26–34=3; 35–44=6; 45–54=3; 55–64=5; 75+=1). The researchers approached 4 additional individuals who subsequently declined to be interviewed.

Coding

The codes for the data analysis were developed via an iterative process. After nearly half the interviews were collected, two of the researchers independently went through the interviews and created an initial set of codes. Following this, the researchers met to discuss the similarities and differences in their initial set of codes and agreed on a codebook. The researchers then used the codebook to code interview data segments via consensus. When appropriate, nested codes and multiple codes were applied to a single segment of interview data. As additional interviews were performed, the researchers reexamined existing codes and made modifications as necessary to the codebook, going back and recoding previously coded interviews. This iterative process was repeated until all interviews were coded and the final codebook was created. All interview responses were coded, regardless of whether or not the interview was truncated.

3.2.2 Interview Excerpts: Participant Snapshots

The next section presents interview results and analysis; however, before I focus on subcomponents of the interviews, I wish to convey a sense of the interviews as a whole. I present below excerpts from three interview sessions. The participants reflect different positions along the spectrum of reactions and different levels of familiarity with AR technologies. The interviews also focus on different underlying themes.

Interview J: The Evolution of Social Norms. Participant J (a 23-year-old unemployed philosophy student and reader/writer), described himself as interested in technology, but did not consider himself a techie. *“I’m straddled between the prehistoric and the modern.”* He was familiar with Google Glass, but did not *“think that their quality was high enough to break into the market yet.”*

J was definitely aware that these kinds of devices can record: *“I would be surprised if their cameras aren’t always on...It would make them easier to interact with, like the*

Kinect for the XBox One. Plus, how else would you fuel the tinfoil hat conspiracy theorists?” On the topic of how he felt about being recorded by such glasses, he said, *“If I got drunk and puked on a friend, I wouldn’t want that out there, but it shouldn’t affect my ability to get elected to public office...There are things that we don’t want in the public, but it won’t be harmful, especially in the future...But, in the interim, people do lose their jobs over Facebook posts.”*

When asked if there are spaces where we shouldn’t record, he replied, *“The extreme example is the bathroom or the bedroom. But it’s only a matter of socialization. Right now it’s not civilized to record in the bathroom. But consider the [Ancient] Greeks. They didn’t use to work out in the nude, until they realized that it was better. So they accepted that.”*

Interviews E & F: Technology and Isolation. Participant E (a 55-year-old female teacher) and Participant F (a 57-year-old male engineer) were interviewed together. E described herself as having *“limited knowledge”* of AR devices, but then proceeded to express an appreciable understanding of the concept. *“A screen comes over the eye,”* she demonstrated, holding up her smartphone to her face, *“and you don’t need a computer; you just cloud WiFi it.”* She was aware that the glasses could take pictures and recordings: *“It seems creepy because they can take pictures surreptitiously. You can go around and take pictures,”* again, she illustrated by using her hand, *“hot girls [click], hot girls [click].”*

While discussing how they would feel about being bystanders to such a device, F chimes in, *“If I really researched privacy issues, I would be more bothered, since it’s probably worse than we know—almost certainly worse than we know...I don’t think the ethical questions have caught up with the technology.”*

E explained, *“I teach young people—18 to 30—and they would probably get the device because it’s the cool new thing. It doesn’t appeal to me. I can’t think of a reason to use them. Technology portrays itself as creating community, but instead it destroys community.”* F added, *“People’s attention spans have been brought down to*

sound bites.”

In response to being asked, both E and F expressed an interest in having AR users ask their permission before taking recordings, but E said, *“I don’t think there’s an actual etiquette for that...or any etiquette for devices in general.”* When asked if they would be interested in technology that would allow them to block themselves from being recorded, both E and F were interested, but F added, *“We probably wouldn’t need to...Once you get to a certain age—over 50—we are invisible anyway.”* E stated, *“In the future technology will let you remove people from videos. I only want to see hot chicks; get rid of people over 25.”*

Interviews V & W: Context and Content Ownership. Participant V (a 20-year-old female dance major) and Participant W (a 21-year-old female dance major) were interviewed together. V had heard of Google Glass, but neither she nor W knew that they could take photos and recordings. W exclaimed, *“Wow, like Spy Kids. It’s real! [laughter]”*

When asked if recording with these types of glasses is similar or different to recording with a cell phone, they expressed that it was different. V said, *“It’s more obvious with a cell phone. It’s like, ‘I’m recording something.’ With the glasses, it’s like, ‘Are you recording my conversation?’ I don’t know. Does it blink?”*

She would find being around an AR device *“a little unsettling—but not too unsettling.”* W elaborated a bit more: *“I’m a dancer, so if I saw a video camera coming down the street I’d probably jump in front of it. [laughs] But if I saw someone coming into a performance—or a movie theatre, I guess—that would be a problem. But if they’re just recording our conversation, it isn’t that interesting.”* Upon being asked about a potential interest in blocking technologies, V explained, *“I’m a broke college student. If it bothered me, I’d approach them. If it got to be an issue—like for working in the theatre—if a lot of people started coming in with these devices I’d probably tell my boss to get one to stop all the recording. That’s actually pretty smart.”*

3.2.3 Results

In this section I present the analysis of interview data. When specific analysis codes appear in the text they are indicated by a bold font. This is a qualitative study that is primarily intended to explore relevant issues. As a result, participant counts should be taken as a rough indicator of the participant population rather than an absolute measure.

Bystander Reactions

At the beginning of each interview, my collaborator or I asked the participant whether or not they had noticed the second researcher's AR glasses. Many of the participants (11/31) had not made any particular note of the glasses (**Noticed Glasses: No**), despite the bias in the sampling methods (see Section 3.2.1).

As the interview proceeded, participants expressed a range of reactions regarding the idea of being a bystander to an AR recording device. As part of the analysis, my collaborator and I coded these sentiments as **AR Bystander: Positive**, **AR Bystander: Negative**, or **AR Bystander: Indifferent**. Participants were split in their reactions, but they primarily either reacted indifferently (16/31) or negatively (12/31) to AR recording; only one participant had a positive reaction. Also notably, some of the participants (6/31) expressed more than one type of sentiment, highlighting the fact that people can have conflicting or complex reactions.

The Familiar: Legality and the Public Stage

When participants offered reasons why recording with an AR device is acceptable or makes no substantive difference to their experience as a bystander, they primarily did so in the context of comparisons with existing technologies. When my collaborator or I probed them, 10 participants indicated that they view AR recording as similar to cell phone recording (**Cell Phone Comparison: Similar**). Some partici-

pants volunteered comparisons to other existing camera technologies. For example, 5 participants specifically commented on the preexisting prevalence of CCTV cameras (**CCTV Comparison**). A few participants made comparisons to other recording technologies, such as the GoPro wearable camera (**Camera Comparison**).

In general, participants rhetorically used these comparisons in one of two ways: to indicate that AR technologies make no difference in the legal landscape, or to indicate that AR technologies make no difference in their expectation of being recorded. For example, Participant N (a 21-year-old male game designer) indicated that he cannot legally stop someone from taking his picture, regardless of device type. Participant AC (a 64-year-old male video producer) is in the Screen Actors Guild; he indicated that no one is allowed to capture his image without written permission. Multiple participants expressed that—between cell phones, CCTVs, and other cameras—they already expect to be recorded whenever they are in public. Not all participants seemed pleased or indifferent about that fact; however, the introduction of AR technologies did not affect their expectations of being recorded. Below are three participant quotes (paraphrased from transcript notes) that illustrate viewpoints along this spectrum:

- Participant L (a 48-year-old female IT manager and informatics student): *I'm fully aware that I'm being photographed all the time. Look at the tracking activities of the police in Boston [referencing the 2013 Boston Marathon bombing]. That was "fantastic," in the literal sense of the word, not necessarily the positive sense.*
- Participant B (a 39-year-old female lawyer): *People are aware that there are a lot of CCTVs around—there's not a street corner in Seattle that's not recorded. It's a bit Big Brother, but we accept it as a society, and it's not like you're in a house.*
- Participant K (a 50+-year-old male who described his occupation as spiritual):

I am consciously sharing just by being present. If I didn't want to be seen I would lock myself up and never go out.

Several participants focused on the concept of appearing in public. This viewpoint is reminiscent of Goffman's theory of the presentation of self in everyday life. In this theory, he describes our interactions as times when we are performing: we are scrutinized by others, and dynamically adapt to their reactions. At other times, we do not wish to be seen, and hide away "backstage" [38].

The Foreign: Subtleness and (Lack of) Prevalence

When probed on the topic, 8 participants indicated that recording with an AR technology is different than recording with a cell phone (**Cell Phone Comparison: Different**). Elaborations on these answers surfaced some reasons why participants regarded these technologies as creating a different experience for the bystander.

Over half (16/31) of the participants—including Participant V quoted in Section 3.2.2—raised the fact that AR glasses are potentially a more subtle form of recording than other form factors (**Subtleness**). Participants indicated that bystanders consequently may not be aware that they are being recorded. This concept of subtleness is somewhat intertwined with the fact that it is relatively easy to initiate a recording (**Ease of Recording**)—an issue that was articulated by 5 participants.

Some participants (8/31) gave another reason why bystanders might not expect to be recorded by AR glasses: the technology's current lack of prevalence (**AR Prevalence**). They indicated that the scarcity of AR devices meant that people would not expect glasses to be recording. In some cases, as in the quote below, the participant explicitly indicated that this expectation would change as the technology becomes more common:

Participant I (a 43-year-old male working in science): *It's slightly more clandestine, but if it gets popular people would be clued in.*

Perspectives on Recording: Who, What, When, Where, Why, and How

Throughout the course of the interviews, participants expressed a number of factors that affected their feelings towards being recorded. For some participants, these factors described why they prefer not to be recorded. Other participants mentioned factors that affect the circumstances in which it is or is not acceptable to be recorded. While some of these issues have been surfaced in prior work (see Section 3.1.1), my collaborators and I show that they arise again: in a different time and place and with a different technology. I present these factors below in approximate order of their prevalence in interview data.

Place. The majority of participants indicated that Place plays a role in whether or not it is acceptable to make recordings. This discussion was predominantly in the context of recording in *public* versus *private*. Some of the participants, however, articulated particular places or types of places in which one should not record. Some of these places were unacceptable by virtue of **Social Norms** (bathrooms, bedrooms, in others' homes). Other locations were described as off-limits owing to existing camera policies (locker rooms, theatres, government buildings, gun stores, some cafés and bars). Participants V and W discuss this issue in their interview (see Section 3.2.2).

Bystander Behavior and Sharing Context. Participants indicated that the acceptability of being recorded was somewhat dependent upon what they were doing at the time (**Bystander Behavior**). For example, one participant did not want an AR user to “shoulder surf” her at the ATM. The majority of the references, however, were in the context of impression management. Again in the context of Goffman [38], we might describe people’s behavior as an interactive performance tailored to a particular audience. When this performance is taken out of context, undesirable or unanticipated consequences can follow. As a result, sharing images or videos online—or the context in which they are shared—affects bystander feelings regarding being recorded (**Sharing Context**). Participant J (see Section 3.2.2) gives

one example of how bystander behavior and sharing can have negative consequences; most participant examples were similar. Participant R (a 35-54 male who works in a mix of entertainment and technology), on the other hand, provides an example where his behavior is not the issue in question, but the sharing context still is: someone else could “superimpose” his recording over a porn film. While this scenario may seem unlikely, it has parallels to media reuse for satirical or damaging purposes.

Perception of Recorder. Participants judged whether or not they minded being recorded based upon their evaluation of the AR user (**Perception of Recorder**). A contextual evaluation is illustrated by the following quote, paraphrased from interview transcripts:

Participant M (a 60+-year-old male retired marine biologist): *I look over at him, I size him up, and if he doesn't look like a pervert—if he just looks like Joe Schmuck—it's not a problem.*

Participants also indicated that the gender of the person wearing the glasses affected their perception of the device.

Other participants expressed evaluating the AR user based upon his or her perceived role; for example, some participants trusted individuals and distrusted corporate and governmental organizations, while other participants had the opposite reaction:

- Participant O (a 32-year-old female dancer, catechist, and graduate student): *Well, he's—I guess he could be from the government or a large corporation—he's an individual, and I feel like that's fairly benign, and I trust that he's not going to do anything too bad with it.*
- Participant L (a 48-year-old female IT manager and informatics student) is concerned about individuals recording, since they are not held to the same moral and ethical bounds as law enforcement.

Identification. Participants articulated concerns about being recorded by AR technologies based on the idea that they—or others—might be identified in the resulting images or videos (**Identification**). Several participants provided further context regarding their concerns:

- Participant A (a 42-year-old female working in customer service) is a foster parent and is concerned that her foster children might be identified in footage.
- Participant Q (a 35-54 male who works in a mix of entertainment and technology) is concerned that he might be tagged in a video alongside a person of interest or a criminal element, resulting in “*guilt by association.*”
- Participant AE (a 43-year-old female who works in social services) is concerned that victims of domestic violence might be identified online, facilitating abusive ex-partners “*coming after them.*”

Vexation. A few of the participants indicated that they would object to being recorded only if it presented an interruption or an irritation (**Vexation**)—if the AR user was “*up in their space*” (Participant B, 39-year-old female lawyer) or “*disturbing*” them (Participant G, 22-year-old male retail worker going to school for graphic design).

Exploring Consent and Control

One of the interview questions probed whether or not participants would want someone to ask them before recording them with AR glasses. The follow-up question asked participants if they would be interested in a device that could block others from recording them. These questions were intended to: (a) surface relevant underlying issues; and (b) explore whether or not a technological mechanism supporting notification, consent (e.g., [10, 14, 77]), or blocking (e.g., [69]) would be of interest to participants.

Permission. Most of the participants (17/31) expressed that they would prefer for someone to ask their permission before recording them with AR glasses (**Permission: Yes**). 7 of them would prefer not be asked or expressed indifference (**Permission: No / Don't Care**). 7 of the responses were uncodable due to ambiguity, truncated interview, or omission (**Permission: Uncodable**). Responses were frequently accompanied with caveats. For example, some participants expressed that they would wish to be asked, but that it is not practical for the AR user to do so (**User Feasibility**). Other participants wished to be asked, but expressed a sense of **Helplessness** regarding their ability to enact their preferences.

For many participants, whether or not they would want their permission sought was dependent upon whether or not they were the focus of the recording (**Focus of Recording**).

Blocking. 12 of the participants expressed an interest in a device that would allow them to block others from recording them (**Blocking: Yes**). 6 of the participants were not interested in such a device (**Blocking: No**). 13 participants' responses were uncodable due to interview truncation, omission, or ambiguity (**Blocking: Uncodable**). Participants variously expressed that their interest was dependent upon: (a) the cost of the device in question; (b) whether or not they would have to wear the device (versus installing an app on their phone); (c) device size; (d) effort involved in using the device; and (f) the prevalence of AR recording devices (for some participants, a prevalence of AR devices would encourage them to use a blocking technology, while for others it was the opposite).

Some participants expressed an interest not particularly for the purpose of blocking AR recording, but for the ability to use them on recording technologies in general:

Participant O (a 32-year-old female dancer, catechist, and graduate student): *ABSOLUTELY. [emphasis in transcription] Not so much for the glasses—I trust the average Joe—but for the cameras everywhere else.*

3.2.4 Discussion

The interviews took place in cafés in Seattle: a city with multiple universities and a concentration of technology companies. My collaborators and I expect that general bystander perspectives regarding recording will shift by city, region, and country. Moreover, while cafés are a rich source for study, they do not capture the full scope of human behaviors. This methodology could be extended to a variety of location types with pertinent theoretical properties, such as: power dynamics (e.g., workplace); specific population types (e.g., playground); disheveled appearances (e.g., gyms); or casual atmosphere (e.g., bars).

In this study, my collaborators and I investigate how individuals perceive AR-style recording in comparison to other classes of recording devices. Participants were split as to whether or not AR devices create a substantively new bystander experience; those who found it different cited subtleness, ease of recording, and the current lack of prevalence as the relevant factors. The scarcity of AR devices is not an inherent property of the technology; however, it can contribute to whether or not an individual expects to be recorded. It remains to be seen whether these factors continue to be perceived as relevant as the novelty of the technology fades.

People frequently: (a) are unable to adequately assess their reactions to a technology before they encounter it—an obstacle which my collaborators and I attempt to lessen with the interview methodology (see Section 3.2.1); (b) change their perceptions with repeated exposure to the technology; or (c) change their views as they become active users of the technology (e.g., [68]). Data gathered now may or may not reflect how individuals will perceive AR recording in the future; either way, it can be used to characterize the adoption of an emerging technology.

Participants expressed interest in the possibility of being asked permission and being able to block recording devices; however, they expressed concerns regarding feasibility and convenience. These factors suggest that privacy-mediating technologies

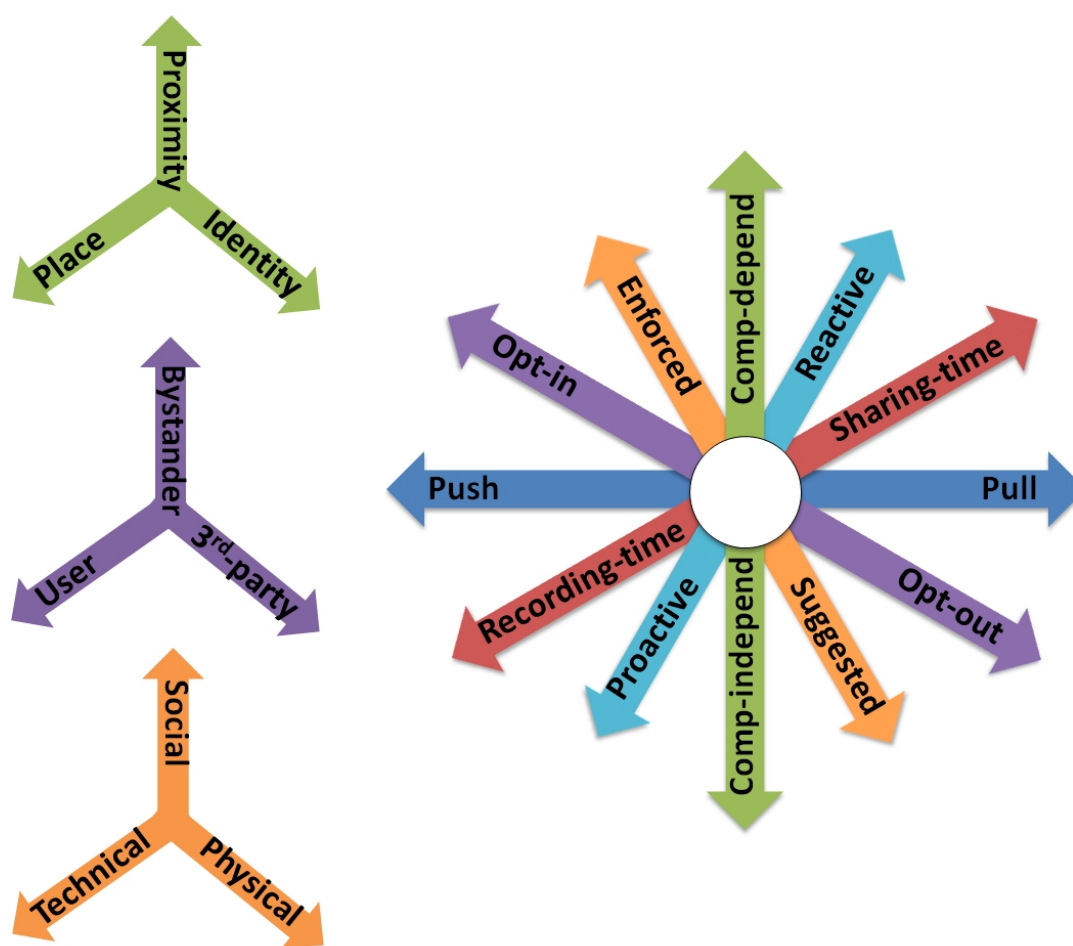


Figure 3.2: A breakdown of potential design axes for privacy-mediating technologies. See Section 3.3 for further discussion.

are a space that merit further research.

3.3 Design Axes for Privacy-Mediating Technologies

The themes that emerge from the results and from considering related literature suggest a number of design considerations, both for AR technologies themselves and for companion technologies in the recording ecosystem. I discuss them in the following subsections and ground them with research references.

Figure 3.2 presents a set of axes by which to characterize or explore the design

space of privacy-mediating technologies; these axes provide a framework with which to consider the following discussions. In alphabetical order, the (largely orthogonal) axes are: (a) **Compliance-dependent** vs. **Compliance-independent**; (b) **Enforced** vs. **Suggested**; (c) **Opt-in** vs. **Opt-out**; (d) **Physical** vs. **Technical** vs. **Social**; (e) **Proactive** vs. **Reactive**; (f) **Proximity-based** vs. **Place-based** vs. **Identity-based**; (g) **Push** vs. **Pull**; (h) **Recording-time** vs. **Sharing-time**; and (i) **User-based** vs. **Bystander-based** vs. **3rd-party-based**. Illustrative references to the axes in the following sections are indicated using an italic font.

Offsetting Subtleness and Negotiating Permission

As noted by participants, one of the key ways that AR technologies are different from other technologies is the subtleness of the recording experience for bystanders. Additionally, participants expressed an interest in being asked permission or being able to block recordings. I explore mechanisms for notification, blocking, and permission below.

Physical Measures. Subtleness may be partially offset by visual or aural cues to bystanders that a recording is taking place. Unfortunately, this method runs the risk of being bypassed by malware (e.g., [16]) or by non-compliant devices. Alternatively, devices could be designed such that their cameras may be physically blocked by switches or shutters.

Technical Measures. The possibility of push-pull interactions leads to an array of potential notification and permission mechanisms. For example, an AR device could push notifications to nearby cell phones that a recording is taking place (*push, user-based*). Such a notification could include information about where the recording might be posted. The AR device could solicit privacy preferences from bystanders' devices (*pull, proximity-based*). Alternatively, bystander cell phones could choose to broadcast their owners' privacy preferences (e.g., [10, 14]). Continuing with this example, the AR user might choose to respect the preferences of bystanders and keep

a recording private (*social, suggested*). The system could also support sending automated notifications to bystanders if relevant photos or videos are posted (*sharing-time*). One way to support these interactions would be to rely upon messages exchanged while the devices are co-located (*proactive*). Alternatively, the system could cryptographically support sending notifications after the fact while still supporting all parties' anonymity (e.g., [59]).

At the other end of the spectrum, privacy preferences could be technically enforced rather than suggested (*compliance-dependent, enforced*). For example, a system could (mostly) guarantee that all bystanders have the ability to take down a recording at any time in the future (e.g., [43]) via a system that is somewhat analogous to digital rights management. Similarly, some individuals—including some of our participants—may have interest in a technology which actively blocks cameras' ability to record them, with or without the operator's cooperation (*bystander-based, compliance-independent*, e.g., [69]).

A discussion of recording preferences and blocking naturally segues to ethical, philosophical, and legal discussions about the ownership of space, the rights of an AR user, and the ownership of content. Many spaces where recording devices are used are privately owned. As such, the owners or event managers might wish to enforce their own policies (*place-based*, e.g., [11]). On the other hand, such a mechanism has the ability to limit individuals' ability to capture and express material. An individual might want to—or have the moral or legal right to—record for a variety of purposes, including: the creation of digital memories for informational or emotional purposes, self-protection, journalism, or social justice. Another question arises once a recording is created: who has ownership over the data? Although there are exceptions, the current model in social media networks is that the content is owned and managed by the uploader. This can create tensions between the media owner and any subjects in the content. Subjects can manually or automatically ask the owner to untag, restrict access to, or remove the content (*suggested, reactive*, e.g., [13]); however, this does

not necessarily circumvent social conflict. Further afield are models for collaborative management of media content (e.g., [84]).

Place as a Social Construct

Previous research has found that the acceptability of recording varies by location, and this study is no exception. Participants indicated that spaces such as homes, locker rooms, and theatres require special treatment. Location and space have definite social and societal meaning, and I do not dispute that there is value in supporting space-based restrictions on recording.

I suggest, however, that designers and technologists consider the broader view of place, rather than space. I use the word “place” to encompass the social characteristics of a space as situated in time and space [46]. For example, an auditorium is a type of space which can at different times host a children’s play, an Alcoholics Anonymous meeting, or a burlesque show; each of these events constitutes a different place and has different accompanying social expectations.

While it may be more difficult to form automated decisions based on a social context than on physical location, it is also a more meaningful distinction. Devices could attempt to gather such context based on co-located individuals, online listings, physical artifacts in the environment, or even the user’s calendar entries. For example, calendar invites or event locations could include recording policies. Even further afield, AR devices which have “prior knowledge” about a given event space’s recording policies could broadcast that knowledge to surrounding devices. How to effectively sense and operate on these data streams remains, however, a non-trivial research problem.

Identification Mitigation

Several participants expressed discomfort with the idea of being recorded on the basis that it facilitates identification. In some cases, participants gave reasons why

identification could lead to negative consequences, including bodily harm. Below I explore some potential ways to mitigate this concern.

Individuals might choose to wear opt-out markers (e.g., [77]) if they do not wish to be recorded. Conversely, they could wear opt-in markers if they do not mind being recorded (*opt-in, compliance-based*). However, outside of specific, structured environments, this strategy is most likely unrealistic.

Counterintuitively, if AR devices could rely upon facial recognition to identify everyone in an image, they could then use that information to blur or obfuscate individuals who have previously expressed or registered that they do not want to be recorded (*opt-out, proactive*). It should be noted, however, that this avenue puts the responsibility of registering on the bystander. Moreover, this approach potentially leaks as much private information as it protects. At the other end of the spectrum, facial recognition of acquaintances could be utilized to anonymize everyone who is not an acquaintance, thereby protecting bystanders.

The above approaches could be used to suggest to AR users that they avoid sharing media with sensitive identifications (*identity-based, suggested*). These approaches could also enforce recording deletion or obfuscation (*enforced, compliance-dependent*). As previously discussed, preventing or altering recordings raises questions of the AR user's rights to aesthetic and accurate memories—not to mention the implications regarding legal evidence (e.g., [17]).

Further afield, bystanders or social media platforms could run independent “watchdog” software (*third-party, reactive*). These agents could review media where the bystander might appear based on metadata such as time and place (e.g., [47]). This approach would allow bystanders to monitor their appearances in public data without relying upon ecosystems or interoperating protocols (*bystander-based, compliance-independent*).

3.4 Summary

Glasses-style AR devices are starting to enter the commercial marketplace. The recording capabilities of these devices have the potential to increase the frequency with which bystanders are recorded in publicly accessible locations. While there has been much controversy in the media surrounding these technologies, little is known about how the general populace perceives such devices.

We sought to help address this knowledge gap with an in-situ qualitative study: my collaborator and I wore a glasses-style AR mockup in cafés and conducted semi-structured interviews with café patrons. Subsequent analysis yielded a variety of information: for example, participants described AR recording as different from other types of recording due to its subtleness and the current scarcity of AR devices. Participants also surfaced factors that make recording less acceptable. For instance, their reactions to recording can be affected by their perception of the AR user and whether or not they can be identified in the recording. Many participants expressed interest in being asked permission or being able to block recording devices; however, they were concerned about the logistics of such capabilities.

I discuss a range of such technologies in Section 3.3. Furthermore, I characterize these systems by supplying axes for design directions (e.g., *proactive* vs. *reactive*, *enforced* vs. *suggested*, *technical* vs. *physical* vs. *social*). The investigation of such technologies is timely: the nascence of AR technologies can potentially be used to bootstrap the inclusion of privacy-mediating measures. The utility of such measures and the utility of characterizing axes extend beyond AR devices to new classes of emerging technologies.

The fact that AR technologies are nascent affords opportunities to the research community. Since these devices are not yet common, we can study how perceptions and usage develop throughout their adoption. Moreover, we have the opportunity to explore privacy-mediating mechanisms; the user experience for AR devices has not

yet become standardized.

The methodology used in this work demonstrates one way to attempt to ground participants' responses to a technology with which they are less familiar, which is an issue that arises frequently when studying emerging technologies. The results of this study—and these kinds of studies in general—help designers and technologists understand people's objections to a technology—whether or not these objections are reflected in the actual technology's functionality. In turn, these objections can be used to drive design decisions to produce technologies that are accepted and that consider the priorities of all stakeholders.

3.5 Acknowledgments

This work was supported in part by the US National Science Foundation (NSF) under awards CNS-0846065. Any opinions, findings, and conclusions or recommendations expressed in this dissertation do not necessarily reflect the views of the funding agencies.

I thank the participants and the study locations for their time. I thank David Molnar for the loan of his equipment. I thank Eun Kyoung Choe, Karen Fisher, Jaeyeon Jung, Julie Kientz, Adam Lerner, David Molnar, Alex Moshchuk, Oriana Riva, Franziska Roesner, Helen Wang, and Information School students for their input. I thank the Pervasive Computing Intel Science and Technology Center. I thank Security Innovation.

Chapter 4

CONTROL-ALT-HACK: A TABLETOP CARD GAME FOR COMPUTER SECURITY AWARENESS AND EDUCATION

My experiences in my research, and particularly my experiences during my work with implantable cardiac devices (Chapter 2) and augmented reality (Chapter 3), have demonstrated to me that there is a growing divide between the realities of computer security and the understanding of users and technologists. In particular, people do not always consider the potential ways that they can be harmed by computer security breaches and all the technologies that are impacted by computer security. As mentioned in Chapter 1, emerging technologies exacerbate this situation. They are capable of sensing new things, they can physically affect the environment around us in new ways, we are relying on them more broadly and deeply in our daily lives, and they are increasingly interconnected. Together, these properties mean that people can be harmed in novel or amplified ways. Moreover, many of these technologies bear little resemblance to desktops or laptops, and may not raise associated security concerns.

I believe that it is necessary to raise people's awareness of the security and privacy issues surrounding emerging technologies. This can be accomplished in a variety of ways, including via more intuitive user interfaces or via educational campaigns. In my research, however, I wished to explore how one could create a tool designed to impart high-level security information implicitly. In particular, I wished to harness the engaging and voluntary properties that are associated with play. My collaborators and I designed, produced, distributed, and evaluated Control-Alt-Hack: a recreational

tabletop card game about computer security. This chapter details that work.

Section 4.1 provides additional overview on the project. Section 4.2 provides more information on the design of the game. Section 4.3 provides details on how I evaluated the game for usage in an educational context. Section 4.4 summarizes this chapter in the context of my dissertation.

Part of the material in this chapter first appeared in [26]. Collaborators on the game design included Tadayoshi Kohno and Adam Shostack. Adam Lerner was a collaborator on the evaluation.

4.1 Overview

Exposing many different kinds of individuals to ideas that make them think about computer security—however briefly—could potentially benefit the status of computer security as whole:

Current and Future Users. The more people prioritize security, the more they might express it with their purchasing power, and the more willing they might be to engage in security and privacy behaviors that require time or effort.

Current and Future Developers. The more developers prioritize security, the more willing they might be to take action. This might mean taking security training, refreshing their knowledge of best security practices, taking more care with their code, or simply thinking to reach out to their institution’s security team.

Current and Future Management. If management prioritizes security, they might dedicate more resources to developing and maintaining secure products and systems, or reward security-promoting behaviors via the institution’s incentive structure.

Future Technologists. I wish to encourage as many people as possible to consider computer security and computer science as a profession, in order to increase the strength of the field as a whole.

There are many avenues to increase people’s awareness of security: publicity cam-

paigns, integration into popular culture, and education and training are just a few. In my work, my desire to create an artifact that exposes people to thinking about security and that facilitates ad hoc, social interactions led my collaborators and I to design *Control-Alt-Hack®: White Hat Hacking for Fun and Profit*: a recreational, tabletop card game about computer security. As of May 2013—the date of the evaluation—approximately 800 requested copies of Control-Alt-Hack had been shipped to 150 educators. (As of July 2013, approximately 3000 copies of Control-Alt-Hack have been given to educators.)

I sent those 150 educators surveys, and 22 educators representing over 450 students submitted feedback about their experiences using Control-Alt-Hack inside and outside of the classroom. Analysis of the evaluation data has indicated that we have had some success meeting our design goals. The contributions of this work are as follows:

An Unconventional Tool. My collaborators and I designed, produced, and distributed of an unconventional tool—specifically a physical game—in order to raise overall awareness and alter perceptions about security.

Evaluation in Context. An evaluation of the card game’s usage in educational contexts.

4.2 Game Design

Goals

Awareness Goals. My primary goal was to increase people’s awareness of computer security needs and challenges, so that they can be more informed technology builders and consumers. This includes:

1. Increasing understanding of the importance of computer security, and the potential risks with inadequate security safeguards.
2. Conveying the breadth of technologies for which computer security is rele-

vant, including not only conventional computing platforms like laptops and Web servers, but also emerging platforms like pervasive technologies and cyber-physical systems.

3. Improving understanding of the diversity of potential threats that security designers must consider and the creativity of attackers.

Perception Goals (Secondary Goal). I additionally seek to show that the information technology community and its professions are open to people of diverse backgrounds. Providing even fictional role-models could help encourage interest in computer science and computer security. More specifically, I aim:

1. To work against negative, dissuasive, or niche stereotypes about people in these fields, and to allow players to identify with one or more of the characters in order to envision themselves in the field.
2. To highlight the variety of professional and personal opportunities available to people with these skills.

Exposure Goal. I seek to have as wide an impact with our Awareness and Perception Goals as possible—the more people that play this game, the more opportunities our game has to increase awareness or change perception.

Why a Game

I believe that games are well positioned to address our specific project goals. If designed well, games can be an appropriate tool for seeding a large audience of people with a modest amount of security information. Briefly:

- Games can be fun, which gets people engaged.
- Games can give you permission to explore ideas and ask questions.

- Games are intended to have intrinsic entertainment value, which gets people to pick them up and use them on their own time.

Given the subject matter, it may seem natural to have created a computer game, rather than a physical tabletop game. Both formats have their merits and their limitations, and in creating this tool my collaborators and I chose to explore the problem space via a physical game. Part of our reasoning in doing so was to take advantage of some of the following factors:

- Physical games may appeal to people who do not enjoy computer games.
- Aside from requiring a surface on which to play, physical games generally do not require extensive setup or have resource dependencies.
- Having a game lying around in a physical space provides the opportunity to read through some of the cards, even if the game is not being actively played.

While the following properties are not exclusive to physical games:

- Physical games can create social environments, which can foster interaction and discussion of ideas encountered.
- Because physical games can create interaction between players, they are suitable for use in social gatherings.

Target Audience

No game strongly appeals to everyone. While my collaborators and I sought to make our game as broadly appealing as possible to raise security awareness within a very large audience, it is most practical to target a specific demographic.

Primary Education Audience. The primary target audience is people with an affinity for computer science and engineering but without requiring significant computer security education, training, or experience. This project targets in particular

those who are early in their careers, including computer science and engineering undergraduate students, high school students, and recent graduates. For example, a high school student in AP Computer Science might play this game, as might a recent hire in software development, test, or management. This goal means that our primary target audience is technically inclined and consists of roughly 15- to 30-year-olds.

Secondary Education Audiences: High school and undergraduate students in the Science, Technology, Engineering, and Math (STEM) disciplines; software developers; gamers; and the broader public. This audience presumably would benefit more from success with the Perception Goal than would the primary audience.

Security Community: As a vector for increased dissemination.

Related Work

I separate our discussion of related work into work on commercial games and games created more as research endeavors.

Commercial Games. Previous commercial tabletop card games dealing with computer security include games such as Fantasy Flight Games' *Android: Netrunner*, published in 2012, and Steve Jackson Game's *Hacker*, published in 1992 (now out of print). I believe that this contribution is distinct in several ways. First, we take many opportunities to ground our card contents in a variety of current technologies and actual attack threats (see Table B.1 in Appendix B.1). While this means that Control-Alt-Hack is at risk of becoming outdated, this also means that it is particularly topical. Second, while these games undoubtedly helped—and continue to help—attract people to computer science and computer security, they portray hacking—and hackers—in the style of a particular niche (although compelling) subculture. My collaborators and I chose, primarily via graphic design and illustration choices, to embrace a more non-traditional hacker “tone” in the hopes of connecting with a slightly different audience. Third, since we created our card game specifically with awareness goals in mind, we are also contributing an evaluation of our game in education contexts.

Games in Research. Gondree et al. [40] gives an overview of some of the benefits of using casual games to impart modest security information; they reference Klopfer et al.'s [53] five freedoms essential to play, and reinterpret those freedoms as mapping to the adversarial, exploratory aspects of computer security.

[d0x3d!] is a tabletop board game designed to casually introduce a wide audience to some of the terminology and adversarial thinking that is involved in network security [1]. *Exploit!* is a card game that is primarily intended for entertainment for the security audience, not education [50]. *Elevation of Privilege* [62, 81], *Protection Poker* [94], and *OWASP Cornucopia* [67] are meant to help train and augment threat modeling and risk assessment in software development.

CyberCIEGE [87] and CyberProtect [88] are electronic games that have players act as network administrators who must utilize limited resources to manage overall network risk.

There are a variety of Capture-the-Flag competitions (e.g., [2, 4]), which are competitive and engaging ways to promote or simulate offensive security. There are also some defensive competitions, such as the Collegiate Cyber Defense Competition [3].

Educational research communities have looked at a variety of aspects of using games in education: for example, making educational games adapt to skill level [7], using game development as a vehicle for programming assignments [78], using games to teach specific topics such as computer ethics [15], or using games to teach how to detect phishing emails [79]. I stress, however, that educational games are used to teach a variety of topics beyond computer science or computer security, such as mathematical fractions (e.g., [32]) or algebra (e.g., [91]).

In the context of security education research, but not in the context of games, there have been numerous explorations of methods for helping students learn the technical skills necessary to protect computer systems against attackers (e.g., [61, 92]).

Choosing Game Mechanics

A game’s “mechanics” includes all numeric and logical elements of the game that contribute to game play; for example, a game’s mechanics might consist of its rules, the number and type of game decks, and the numbers or gameplay actions on those cards. It can be challenging to design mechanics that lead to well-balanced games. Variables include: the number of players; the time it takes to learn the rules; the time it takes to play; replay value; cooperative versus competitive paradigms; the ability to rebound from a losing streak; and the variety of winning strategies. The story, flavor text, and art rest on top of the mechanics.

I initially explored creating game mechanics from scratch. However, since neither myself nor my collaborators are experts in game mechanics, we chose to license a system from a pre-existing game and then create all new game content. This approach allowed us to forgo playtesting the mechanics—a necessary, time-consuming step to ensure game balance and enjoyment. We did do playtesting to review our game content, which we discuss in Section 4.2.

I explored the rules and mechanics of a number of games available for sale in gaming stores for a game that would support our desired design goals. For example, we wanted a game where a player took on the role of a character, so that they could identify with someone in the computer security field (Perception Goals); we immediately gravitated towards games whose characters featured a variety of skills, in order to highlight the somewhat eclectic specializations that can help improve—or break—a system’s security. I also wanted a game that would naturally support a variety of textually-heavy scenarios or encounters.

My collaborators and I licensed the *Ninja Burger* mechanic from Steve Jackson Games [37], best known for their *Munchkin* card game and the *GURPS* roleplaying system. *Ninja Burger* met our above criteria, and we transformed the game into *Control-Alt-Hack: White Hat Hacking for Fun and Profit*. Instead of delivering burg-

ers in fun scenarios in the quest to become the next branch manager, our players tackle a range of technically-themed scenarios with the goal of becoming the next company CEO.

Feedback Process

My collaborators and I solicited feedback on iterations of the Control-Alt-Hack card deck to gather suggestions to improve the game and assess its ability to meet our goals. These formative evaluations took the form of playtest sessions or “show and tell” sessions, and were conducted with a variety of parties, including: undergraduates in an introductory computer science course (n=10); undergraduates involved in a computer security competition (n=5); graduate students affiliated with a computer security lab (n=8); graduate students (unaffiliated with a security lab) who have an interest in gaming (n=2); computer science professors (n=2); a computer science lecturer (n=1); a former high school teacher of computer science, now an undergraduate lecturer (n=1); outreach officers (n=3); and assorted non-experts (n=14). In response to this evaluation feedback, we: changed specific card text, modified art, and added new cards to help keep track of gameplay decisions.

Brief Overview of Control-Alt-Hack

The following is the premise of the game:

You and your fellow players work for Hackers, Inc.: a small, elite computer security company of ethical (a.k.a., white hat) hackers who perform security audits and provide consultation services. Their motto? “You Pay Us to Hack You.”

Your job is centered around Missions—tasks that require you to apply your hacker skills (and a bit of luck) in order to succeed. Use your Social Engineering and Network Ninja skills to break the region’s power grid, or



Figure 4.1: A photo of the game box and contents. Photo Credit: Juliet Fiss. ©University of Washington

apply a bit of Hardware Hacking and Software Wizardry to convert your robotic vacuum cleaner into an interactive pet toy...no two jobs are the same. So pick up the dice, and get hacking!

Figure 4.1 shows the game box and contents. Figure 4.2 shows some of the game art, and Figure 4.3 shows some of the card contents.

Each turn each player attempts a single Mission, so players get to see a number of Missions throughout the course of the game. By incorporating a large number of technologies and security threats into the Mission narratives, the game communicates a variety of security ideas throughout its duration.

Juggling Design Constraints

The game creation process was driven by goals and constraints, some occasionally in direct conflict; seeking optimal solutions (or pleasing compromises) took significant effort and iteration.

In creating the cards' textual content, my collaborators and I balanced a number of goals and restrictions: (1) Including Technical Content; (2) Mapping Game Mechanics; (3) Offering Comprehensibility; (4) Maintaining Brevity; and (5) Incorporating



Figure 4.2: The character art from the portrait side of 12 of the game's 16 Hacker cards. ©University of Washington

Humor.

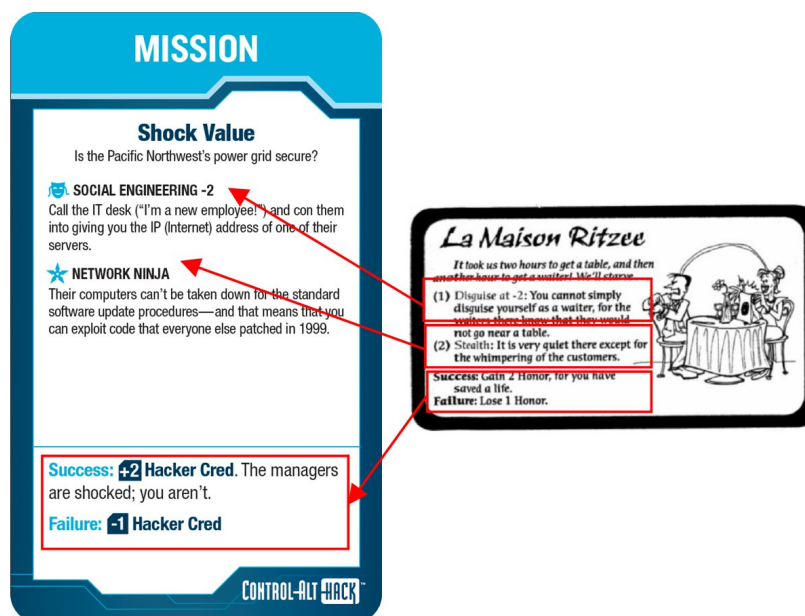
Including Technical Content. My collaborators and I began by creating a list of the content we wanted to cover in order to address our Awareness Goals and convey the range and depth of computer security issues: we brainstormed lists of technologies, attacks, defenses, attacker types, and the range of human assets that can be impacted by system breaches. Table B.1 in Appendix B.1 gives some sample card titles and topics, along with examples of specific research that inspired their inclusion. We sought topics that would be relevant and interesting to players through personal (e.g., social networks), educational (e.g., browser cookies), or professional experience (e.g., patching) or through the news and media (e.g., credit card theft). During the Feedback Process (Section 4.2), we solicited feedback on the selection and technical accuracy of the content which we portrayed.

Most of the game relates to computer security: of the 56 Mission cards, 44 deal directly with security topics, 6 with technological activities (as in Figure 4.3a), and



(a) a Mission demonstrating the usage of technical skills for artistic purposes (Perception Goal (2)).

(b) a Bag of Tricks card illustrating a particular attack threat (Awareness Goal (3)).



(c) a Mission describing a social engineering attack on an SCADA system (Awareness Goals (2) & (3)), along with the mappings to the original Ninja Burger card.

Figure 4.3: Sample card content from Control-Alt-Hack. ©University of Washington

the remaining 6 deal with related topics like puzzles, the role of computer security in history, or the value of professional networking. For content balance and enjoyability, we intentionally did not want all of the cards to focus on computer security topics.

Mapping Game Mechanics. The characters, their skills, and the Missions—which require the use of various combinations of those skills—in Control-Alt-Hack are isomorphic to those in Ninja Burger in order to preserve game balance. Significant iteration and exploration was required to create reasonably realistic and fun story justifications for the combinations of skills required for all 56 Missions. See the Mission “Shock Value” in Figure 4.3c for an example where it was necessary to invent an attack requiring Social Engineering and Network Ninja skills, along with a mapping from the original Ninja Burger card. Similar effort was required to create content for the game’s 72 Entropy cards and 16 Hacker cards.

Offering Comprehensibility. Given our target audience (Section 4.2), our goal of creating enthusiasm for computer security and computer science (Perception Goals), and our desire to reach a broad audience (Exposure Goal), we needed to make our text understandable to those without extensive security experience—without sacrificing technical integrity. My collaborators and I attempted to always make the meaning of terms implicitly clear, explicitly clear, or irrelevant to understanding the overall gist of the card. For example, “Shock Value” in Figure 4.3c parenthetically defines an IP address as an Internet address, and “Dumpster Diving” in Figure 4.3b defines dumpster diving within the text of the card.

Observe how this latter card also incorporates additional learning content: the card helps illustrate that defensive measures (guards, in this case) are not always effective, and that the creativity of attackers can be surprising (such as renting a garbage truck, which many people may not have thought possible).

Incorporating Humor. I incorporated humor into the game in order to make it more enjoyable. The humor primarily (but not exclusively) took the form of: (1) puns; (2) popular culture references; or (3) sexual innuendo, although we attempted

to keep the innuendo tasteful and respectful, and we evaluated the cards with stakeholder groups prior to finalizing them (Section 4.2). For example, “Shock Value” in Figure 4.3c has puns, and “eTextiles” in Figure 4.3a has a popular culture reference in its Hardware Hacking task.

Visuals. My collaborators and I directed illustration and graphic design as part of the game’s content creation process. We purposefully allocated a non-trivial portion of our resources to these visuals for two reasons: (a) to make it easier for players to identify with and project themselves onto Hacker characters (Perception Goals); and (b) to make the game visually appealing, hopefully attracting players (Exposure Goal) and implicitly showing that a focus on technology does not preclude placing importance on aesthetics (Perception Goals).

In creating Hacker portraits, we addressed the Perception and Exposure Goals by balancing the characters’ genders and ethnicities and by showing them engaging in a variety of hobbies. Figure 4.2 shows the character art from the portrait side of 12 of the 16 Hacker cards.

Distribution, Exposure, and Preliminary Impact

In order to reach a diverse set of audiences (Exposure Goal), we chose to make Control-Alt-Hack available via two different avenues:

1. Available for free to educators who submit a request via <http://www.controlalthishack.com>. As of Summer 2013, the supply (approximately 3000 copies) was depleted.
2. Available for sale on Amazon.com via RGB Hats, LLC, which was founded by two of the co-authors and which licensed the game from the University of Washington. This distribution method also allows production of the game to be self-sustaining.

From when the game was made available in November through March, we shipped approximately 800 copies of the game to 150 different educators who requested copies. Approximately 50 copies were also handed out at the SIGCSE 2013 poster session. Together, these educators served as the recruitment pool for our summative evaluation of the game (Section 4.3). Additionally, over 300 copies have been distributed at a variety of NSF-sponsored job fairs, competitions, and similar events.

My collaborators and I were invited to present a talk on the game at a large web company's internal security training conference, and an optional play session was held at the conclusion of the hands-on training.

4.3 Game Evaluation

4.3.1 Study Design

In this chapter I present evaluations of Control-Alt-Hack via two methods:

- **Primary:** Feedback surveys from educators who requested copies of the game.
- **Secondary:** User studies performed with the game. Both methods were approved by the University of Washington's Human Subjects Institutional Review Board.

Educator Feedback Surveys

I distributed online feedback surveys via email to the 150 instructors who received educator copies prior to May 2013. Appendix B.2 shows the questions asked on the educator survey. 22 educators submitted responses to the surveys.

Coding. Two researchers analyzed the survey responses independently and formed preliminary opinions about the categories that emerged from the data. The researchers then compared the categories and formed a cohesive coding scheme via

Responding Participant	Course	Class Size	Student Level	Prior Security Experience	Would have covered [the security material in Control-Alt-Hack] otherwise?	Time Taken	Supplementary assignment involving Control-Alt-Hack
E1-classroom	Information Software Technology	30	HS	No / Some Informal	Yes	60 min	No
E4-classroom	Unknown	12	UG	No / Some Informal	Yes	50 min	Yes
E6-classroom	Computer Science	75	HS	Some Informal	No	75 min	No
E7-classroom	Cyber-Security and Information Assurance	56	UG	No / Some Informal	Yes	120 min	Yes
E8-classroom	Computer and Network Security	10	UG, G	Some Informal / Prior Educational	Yes	120 min	No
E9-classroom	Computers and Information Technology	60	HS	Prior Educational	No	75 min	No
E10-classroom	Game Design	65	HS	No / Some Informal	Yes	90 min	Yes
E12-classroom	Computer Security	22	UG	Prior Educational	No ¹	80 min	Yes
E13-classroom	IT Security	8	UG	Prior Educational	Yes	45 min	No
E14-classroom	Information Security	15	UG	Some Informal / Prior Educational	Yes	120 min	No
E16-classroom	Intro CS Web Design	35	HS	No	Yes	40 min	No
E17-classroom	Cyber Security	2	HS	Prior Educational	Yes	30 min	No
E18-classroom	Fundamentals of Information Security	30	UG	No / Some Informal / Prior Educational / Prior Professional	Yes	75 min	Yes
E19-classroom	Computer and Network Security	27	UG	No / Some Informal	Yes	60 min	No

Table 4.1: Classroom-based educator activity contexts. The shaded cells represent cases of interest, some of which are discussed in Sections 4.3.6 and B.4. HS = high school; UG = undergraduate; G = graduate.

consensus. The primary coder recoded the educator surveys according to this coding scheme. (Complying with the University of Washington's institution's conflict man-

Responding Participant	Context	Time Taken
E2-ACM	Extra-curricular activity with undergrads in the ACM	150 min
E3-vetting	University instructors vetting the game	150 min
E5-no-play ²	Instructor vetting the game with adult friends	N/A
E11-checkout	Provided as a checkout for students to play with friend and family	150 min
E15-vetting	Instructor vetting with graduate students, faculty, and staff	60 min
E20-vetting-didnt-read ³	Instructor vetting	N/A
E21-lunch	Departmental staff lunch	60 min
E22-vetting	Instructor vetting	90 min

Table 4.2: Non-classroom-based educator activity contexts.

agement plan, one of the researchers has no financial interest in RGB Hats, LLC.) In the evaluation, the survey in its entirety was used as the unit of analysis, rather than individual responses; that is to say, if part of an educator’s response received the code “Awareness,” it did not matter which question on the survey elicited the relevant response, and it did not matter how many times the survey was coded for “Awareness.”

The primary coder and the secondary coder had 93% agreement across all educator surveys (N=22) and codes (N=7); there were 11 cases where the primary and the reliability coder disagreed. All cases are provided in Appendix B.3, along with contextual quotes. Except for one case in which the reliability coder misread the data and coded an error, the primary coder’s results—the results reported in the chapter—always represent the stricter of the two viewpoints. That is, these results report the upper bound on the two interpretations of the critiques to the game and the lower bound on the game’s role in engagement and awareness.

The primary and secondary coders independently labeled educator activities as classroom-based or non-classroom-based activities; they had 100% agreement. Table 4.1 and Table 4.2 list information about classroom- and non-classroom-based activities, respectively.

User Studies

I posted recruitment ads inviting participants to join us for a games study session on: an institution-wide electronic bulletin board; and in the local Craigslist gigs listings. I held two game study sessions: one with 7 people (M=3, F=4) divided into two gameplay groups; and one with 4 people in one gameplay group (M=1, F=3). The participants covered a range of ages (mean=31, min=18, max=50, median=29). 5 of the participants could be categorized as “hobbyist” gamers, and 6 had casual or little gaming experience. Each session lasted approximately 2 hours. Participants were compensated \$20 for their time. Following consent paperwork, participants filled out a short pre-gameplay survey. After this participants were shown a 15-minute video introducing them to gameplay; I used a video for consistency between sessions. Participants played for 40–60 minutes, then filled out a short post-gameplay questionnaire.

Coding. Two researchers independently analyzed the survey responses for themes and categories. (Complying with the University of Washington’s institution’s conflict management plan, one of the researchers has no financial interest in RGB Hats, LLC.) The researchers discussed and came to a consensus regarding the data of interest in the survey responses. The data in question is presented as direct quotations, and the goal mappings were decided via consensus coding.

4.3.2 Results

The 22 educators who responded to the survey used the game with over 450 students at the high school, undergraduate, and graduate levels in computer science, computer security, and game design courses. These courses were primarily, although not exclusively, based in the United States. The educator survey results are the primary evaluation of Control-Alt-Hack in this study.

As previously mentioned, the educator survey responses fell into one of two cate-

	Positive Functions		Critiques				
	Social / Engage- ment	Awareness	Takes a long time to learn	Takes a long time to play	Not enough fun	Not enough educa- tional value	Has inap- propriate content
E1-classroom							
E4-classroom							
E6-classroom							
E7-classroom							
E8-classroom							
E9-classroom							
E10-classroom							
E12-classroom							
E13-classroom							
E14-classroom							
E16-classroom							
E17-classroom							
E18-classroom							
E19-classroom							

Table 4.3: Classroom-based educator survey analysis results.

gories: feedback about an activity using Control-Alt-Hack that took place in a classroom, or feedback about an activity using Control-Alt-Hack that took place outside of a classroom. Many of the reported non-classroom activities were from educators who were vetting the game for classroom use, and subsequently decided not to use the game; the other non-classroom activities were an ACM gathering, a lunch activity, and offering the game to students to check out and take home. Table 4.1 provides additional information on the classroom activities (N=14), and Table 4.2 provides additional information on the non-classroom activities (N=8).

4.3.3 Positive Functions

Appreciation of the game expressed in educator surveys generally described the game as fulfilling one of two functions: being fun or serving a social function (Social/Engagement); or increasing students' awareness of computer security or computer science issues (Awareness).

Social/Engagement (Classroom: 11/14; Non-Classroom: 2/8). “Social/Engagement” was marked when the educator was deemed to be indicating that usage of the game was fun, engaging, and/or contributed to serve a social function, such as an icebreaker or a breather before a test. The following quotes are two examples:

- **E7-classroom (56 undergraduates, Cyber-Security and Information Awareness):** *“It worked as a way to break the ice and get students from diverse majors get to know [sic] each other and get thinking about the topics of the course.”*
- **E19-classroom (27 undergraduates, Computer and Network Security):** *“I just wanted to reiterate how great my students thought the game was! The students begged me to leave the game in the student lounge so they could continue to play, and from what I hear it’s made a trip or two out to our weekly majors night at the pub.”*

Awareness (Classroom: 11/14; Non-Classroom: 1/8). “Awareness” was marked when the educator was deemed to be indicating that usage of the game served to increase students’ awareness of security in some fashion, such as: increasing exposure to domain terminology; raising awareness of career opportunities; stimulating discussion; or stimulating critical thinking. The following quotes are two examples:

- **E9-classroom (60 high school students, Computers and Information Technology):** *“The game did not necessarily teach security methods, but it did a great job of teaching vocabulary and literacy.” “It increased awareness of my program, and it got more students interested in computer science.”*
- **E19-classroom (27 undergraduates, Computer and Network Security):** *“They really got into it and there was a lot of strategizing” ... “They were*

mainly focused on causing pain to their classmates, but as I wandered around the room I heard some great discussions about the tradeoffs of choosing various hackers' skill sets, what various missions meant, etc."

Table 4.3 shows the Positive Functions results from the classroom-based educator responses, and Table 4.4 shows the results from the non-classroom-based educator responses.

Overall, in the classroom contexts, 11 of the 14 educators indicated that the game served a Social/Engagement role, and a different set of 11 educators indicated that the game served to increase Awareness. For the educators who did not provide responses that indicated that the game raised awareness, two were courses about computer security; these educators also indicated that the game did not have enough educational content (Section 4.3.4). This suggests that although the design goals were aligned with the intentions of educators not already teaching a computer security course, the goals were not well aligned with some educators' intentions in using the game in security-focused courses.

In the non-classroom contexts, 2 of the educators' responses indicated that the game filled a Social/Engagement role (E2-ACM, E15-vetting), and 1 of the educators indicated that the game helped increase Awareness (E2-ACM). The relative lack of educators reporting positive game functions in non-classroom activities could be a result of the fact that many of the responses in the non-classroom context were from educators who played the game (or not, in 2 cases) out of the classroom in order to vet it for its suitability for use in the classroom. In many of those cases, the educator decided not to use Control-Alt-Hack in the classroom (Section 4.3.5 and Table 4.5), so it is not surprising that they do not comment that the game serves positive functions.

	Positive Functions		Critiques				
	Social / Engagement	Awareness	Takes a long time to learn	Takes a long time to play	Not enough fun	Not enough educational value	Has inappropriate content
E2-ACM							
E3-vetting							
E5-no-play ⁴							
E11-checkout							
E15-vetting							
E20-vetting-didn't-read ⁵							
E21-lunch							
E22-vetting							

Table 4.4: Non-classroom-based educator survey analysis results.

Discussion

Overall, I find that the feedback on the game—in the classrooms in which it was used—shows promising indications that it performs multiple positive functions.

Awareness. In most of the surveys, educators' comments indicated that the game helped raise students' awareness of issues related to computer security. Raising individuals' awareness of the risks, challenges, technologies, and professions involved in computer security was a large part of the purpose in creating the game (Goals, Section 4.2).

Social/Engagement. Again in most of the surveys, educators' comments indicated that the game served a Social/Engagement role in the classroom. This is promising for two reasons: first, it is somewhat correlated with “fun,” which can increase engagement or encourage people outside of the classroom to pick up the game. Second, some of the educators used the game specifically because they had need of a non-traditional educational activity; Section 4.3.6 explores the cases where the educators used the game, but would not have otherwise covered comparable security material. The apparent success of the game's Social/Engagement function, as

represented in the evaluation, suggests that the produced game is aligned with the Exposure Goal.

4.3.4 Critiques and Tradeoffs

The critiques of the game contained in educators' responses were analyzed as falling into one or more of five somewhat self-explanatory statements: (1) Takes a long time to learn; (2) Takes a long time to play; (3) Not enough fun; (4) Not enough educational value; and (5) Has inappropriate content. I discuss the critiques at some length because many of them directly reflect the design tradeoffs that my collaborators and I embraced to meet the intended goals.

Takes a long time to learn. Examples:

- **E5-no-play:** *“Honestly, after reading over the rules, we didn’t understand how to play it, and we gave up. So sorry!”*
- **E15-vetting:** *“The game itself is too complex to easily teach and use for the first time.”*

Following the shipment of the game to educators, I have created a new video that walks viewers through game setup and gameplay in a shorter, clearer format (a video of an hour-long conference talk was previously available which contained an explanation of how to play), which I will publish online; the new video is 10 minutes long. Some of the learning curve is due to the complexity of the game mechanics that we chose (Section 4.2); however, we accepted a level of complexity as a good tradeoff for increased replay value and the in-game opportunities to strategize.

Takes a long time to play. Examples:

- **E14-classroom (15 undergraduates, Information Security):** *“Shorten the game and eliminate some components.”*

- **E12-classroom (22 undergraduates, Computer Security):** *“Students reported that they enjoyed the game, but that the hour twenty was pushing the limit.”*

Gameplay duration can vary depending upon the number of players, players’ familiarity with the rules, and the emergent characteristics of a particular game instance. Potentially long gameplay can make the game unwieldy for the classroom setting; however, the gameplay duration can be an asset in other social settings. Many educators indicated positive results even when playing a version of the game truncated to fit into a class period.

Not enough fun. One example (the only other instance an example of coder disagreement, and is given in Appendix B.3):

- **E3-vetting:** *“The feedback from the instructors trying the game is that it didn’t seem very enjoyable to play or strategic. It may be that more experience will change this, but the first impression was not positive.”*

While the players in the above example (adult instructors) are not the primary target audience, there is no guarantee that the instructors’ students would have found the game fun. I do not have sufficient data to confidently predict who will or will not enjoy the game; nevertheless, observation and anecdotes suggest at the very least that if the audience is familiar with and enjoys the style of game on which Control-Alt-Hack is based, then it is relatively likely that they will find the game fun.

Not enough educational value. Examples:

- **E11-checkout:** *“The game could use more specificity around computer activity. My students were hoping for a higher level of rigor.”*
- **E17-classroom:** *“Since we approached the game expecting to be tested on our knowledge of vulnerabilities and penetration techniques, we were dissatisfied in that manner, but we enjoyed the overall concept.”*

My collaborators and I intentionally chose a lower level of technical depth in the design phase in order to further the Exposure Goal and be comprehensible to a wider portion of the target audience; in the case of these classrooms that decision was not well aligned with instructors' intentions. I recognize that the game is not a good fit for students with a more advanced security background who are hoping to learn new material; this would only be accomplished if the game were paired with a supplementary activity, as some educators chose to do (see Appendix B.4).

Has inappropriate content. There is only one instance of this critique appearing in the data:

- **E20-vetting-didn't-read:** *"I didn't have time to vet the game for appropriateness and, from what I did read on the above site, I felt that the cards significantly contributed to a learning environment hostile to women."*

My collaborators and I do not wish to create an environment hostile to women, and kept gender issues at the forefront of our minds during game development. We took care to make references gender-neutral or gender-balanced: for example, the CEO is a woman, half of the Hacker cards are female, and with one exception, all innuendo is gender-neutral (a Mission card about cell phone security has the title "That's What She Said"). We recognize that innuendo can make an environment more hostile to women, particularly if the environment already has uncomfortable overtones; however, during the design phase we gathered feedback on the appropriateness of the content from multiple parties, including a former (female) teacher of high school computer science and 3 (female) outreach officers (Section 4.2), and incorporated it into the game. For example, we redid the style of dress of one of the female Hackers in response to their comments. The materials we distribute to educators included a list of PG-14 cards which can be reviewed for content and/or removed from the deck.

Discussion

Table 4.3 presents the classroom-based educator experiences coded for the goals and critiques; Table 4.4 presents the same for the non-classroom-based educator experiences.

The most prominent critique was that the game takes a long time to learn (classroom: 4/14, non-classroom: 3/8). From observation and anecdotes, individuals who are familiar with this style of game find it fairly quick to pick up. For example, E13-classroom gave this quote: *“The students with some game experience found it obvious and intuitive. They would say “this is easy.”*” Additionally, we suggest that educators could make the start of gameplay smoother by pre-designating individuals to learn the rules and play together ahead of time, so that those individuals can then seed gameplay groups during the activity.

The second most prominent critique was that the game did not have enough educational value (classroom: 4/14, non-classroom: 5/8). As previously mentioned, many of the non-classroom educators reported on the experience wherein they vetted the game, and chose not to use it in their classroom. Control-Alt-Hack may not be suitable for all educational contexts, but its educational value can be increased by pairing it with or using it to bootstrap a level-appropriate supplementary activity, as done by 5 of the classroom educators.

The third most common critique—and the only other critique expressed by educators who used the game in the classroom—was that the game took too long to play. Control-Alt-Hack may not be suitable for all class formats and in all contexts; however, from observation and anecdotes we tentatively find that having more than 4 players in a game significantly extends the duration of gameplay; we therefore suggest staying below 5 players in a game. Responses indicate that there is some value in playing a short game, even if players do not have time to finish; educators who provided as little as 40 minutes of time to play (E16-classroom) reported some positive

results. Additionally, gameplay is somewhat modular, with logical periodic stopping points; if players are already familiar with gameplay, then individual rounds are of manageable lengths.

4.3.5 “Would Use Again”

Educator	Would Use Again	Would Suggest to Others
E1-classroom	Yes	Yes
E4-classroom	Yes	Yes
E6-classroom	No	Yes
E7-classroom	Yes	Yes
E8-classroom	No ⁶	Yes
E9-classroom	Yes	Yes
E10-classroom	Yes	Yes
E12-classroom	No ⁷	Yes
E13-classroom	Yes	Yes
E14-classroom	Yes	Yes
E16-classroom	Yes	Yes
E17-classroom	No	No
E18-classroom	Yes	Yes
E19-classroom	Yes	Yes
E2-ACM	Yes	Yes
E3-vetting	No	No
E5-no-play	No	No
E11-checkout	Yes	Yes
E15-vetting	No	No
E20-vetting-didn't-read	No	No
E21-lunch	Yes ⁸	Yes
E22-vetting	No	No

Table 4.5: Classroom-use and non-classroom-use responses as to whether or not educator would use the game again, and whether or not the educator would suggest the game to others.

To serve as an overall assessment of the game’s usefulness, we asked educators the following questions on the surveys:

Would you use Control-Alt-Hack again in your classroom? Why or why not? Would you suggest Control-Alt-Hack to others? Why or why

not?

Educators' responses are given in Table 4.5. Overall, the results are promising. 13 of the 14 educators who used the game in their classrooms reported that they would suggest the game to others, and 10 of them reported that they would use Control-Alt-Hack again. E8-classroom responded that they would not use the game with those who already had some familiarity with the subject, but might with high school students or interns, and E12-classroom clarified that they would not use the game again in class due to time constraints, but might as an out-of-class exercise; for both of these educators, this suggests that they still find merit in the game, even if it is not an appropriate match for their instructional needs. E17-classroom indicated elsewhere in responses that the game did not contain sufficient educational content (Table 4.3), so we surmise that is why they will not use the game again or recommend it to others. As mentioned in the previous section, the educational level of the game was an intentional decision related to the Primary Audience and Exposure Goals (Section 4.2).

For the non-classroom experiences with the game, 5 of the educators were playing the game with other instructors, friends, graduate students, or staff to vet its use, and did not subsequently report on using the game in their classrooms. The remaining 3 educators would use the game again and would suggest it to others (E2-ACM, E11-checkout, and E21-lunch). E21-lunch clarified that they might use the game again and recommend it to others, but only with supplementary educational material and after further consideration. These three scenarios—an extracurricular club, a checkout, and a staff lunch—are highly aligned with the social, ad hoc interaction model supported by choosing to create a recreational game.

4.3.6 Reaching New Audiences

Interestingly, 2 of the 14 educators who used Control-Alt-Hack in their classrooms reported that they would not have covered similar security material in any other format. An additional educator gave this response, but was teaching a computer security course (E12-classroom), so this response may have been in error or a misinterpretation of our intention when posing the question (*If you had not used Control-Alt-Hack, would you still have covered the material?*); it is also possible that the educator intended to convey that they would not have covered topics included in the game such as physical security or cyber-physical security. The results are conservative and count this response as an error. If the educators in the remaining contexts would not have covered comparable security material, however, then these classrooms represent instances where the game can serve to increase security awareness, presumably precisely because of its non-traditional format:

- **E6-classroom:** 75 high school students in a Computer Science course with some prior informal security experience.
- **E9-classroom:** 60 high school students in a Computers and Information Technology course with prior educational security experience.

This exposure of individuals in the Primary Audience (Section 4.2) to more security content than they might otherwise have been exposed is an indication of success.

4.3.7 User Study Results

With the educator surveys—the primary evaluation method of Control-Alt-Hack in this study—we gained the valuable perspectives of informed and expert individu-

⁹This was actually in response to the question: *Now that you've performed the activity, what do you think of when you think of computer security? (This may or may not have changed.)*

¹⁰These goals are only potentially implicated in the response. I invite the reader to perform personal interpretations.

Participant	Participant Quote	Goal Mappings
A	“Slightly. I was aware that active testing and debugging are needed to improve security + add to innovation, but the reminder was helpful. The game led me to think about some aspects of modern life I don’t usually consider.”	Awareness #2: Breadth of Technologies
B	“I have to be honest and say that I’ve never heard of a “white hat” hacker before. I’ve always associated hackers with a negative term. Computer security consists of a lot more tasks than I had at first thought it had. Computer security applies to a lot of areas, like cars and phone apps, which I hadn’t thought of.” ⁹	Awareness #2: Breadth of Technologies
C	—	
D	—	
E	“Not much. There was stuff such as not leaving laptops or usb drives out where others can get at them that I had known about but never gave much thought to before.”	Awareness #3: Creativity of Adversaries
F	“Little bit w/ thinking of different scenarios like the small level computer hacking. In general I think of bigger hacking crimes when I think of hacking.”	Awareness #1: Importance/Impact of Security ¹⁰ Awareness #2: Breadth of Technologies ¹⁰
G	“No except that hacking might be fun to use the knowledge to help solve a problem.”	Perception #2: Professional Opportunities ¹⁰
L	—	
M	“Yes. I didn’t give much thought to it before or how many different ways it could be approached.”	Awareness #3: Creativity of Adversaries ¹⁰
N	“No, except that its [sic] very complicated.”	Awareness #3: Creativity of Adversaries ¹⁰
O	“Certainly lightens the mood for my outlook on C.S. and sheds some light for understanding reality of tasks involved.”	Perception #1: Counter-Stereotype Awareness #3: Creativity of Adversaries ¹⁰

Table 4.6: User responses and mappings to the design goals. Participants with no quotes did not provide evidence indicating that their awareness or perception of computer security changed. Project goals are fully articulated in Section 4.2.

als, as well as secondhand access to a large population of students. I also, however, wished to more directly study individuals’ experiences with the game, and therefore performed a supplementary user study. The sessions primarily simulated the experience of individuals of varying backgrounds picking up and playing the game in a non-classroom-setting. Section 4.3.1 provides background on the participants.

In performing the user study, we received participant responses that indicated that—at least in the short term—we are increasing or reinforcing participants’ aware-

ness and/or improving their perception of computer security and computer science, as per the Awareness and Perception Goals articulated in Section 4.2.

Table 4.6 presents participant quotes in response to the prompt on the post-gameplay questionnaire:

After performing the activity, some people say that their perception of computer security has changed, while others don't feel that it has changed much at all. Would you say that your perception of computer security has changed? If so, how?

8 of the 11 participants provided responses which gave some indication that their awareness of computer security issues increased or their perceptions about the field were changed. Interestingly, even though some of these participants responded that their perception of computer security had not changed (2 out of the 8), they proceeded to elaborate and provide qualitative evidence that they were engaged with one of the learning goals. For the remaining 3 participants, none of their responses suggested that their awareness had increased or that their perceptions had changed. Some participants (3/11) supplied critiques on the game; however, the sentiments in those critiques are covered by the educators' critiques (Section 4.3.4), and we do not discuss them further here.

There is a range of participant responses present even in the small sample size. The Goal Mappings column provides a loose mapping from the participant's response to the project goals; the process is subjective, and we invite readers to interpret different mappings from participant responses to project goals.

The project goals are fully articulated in Section 4.2, but they might be paraphrased and shortened as follows:

- Awareness Goal #1: Importance/Impact of Security;
- Awareness Goal #2: Breadth of Technologies;

- Awareness Goal #3: Creativity of Adversaries;
- Perception Goal #1: Counter-Stereotype; and
- Perception Goal #2: Professional Opportunities.

All of the goals appear at least once in Table 4.6, suggesting that we have had some success in crafting the game to touch upon the issues in question.

4.3.8 Discussion

I take this opportunity to discuss some of the reflections from my collaborators and I from going through the process of creating, distributing, and evaluating a computer security-themed tabletop card game for the purpose of promoting computer security awareness and education.

Physical Games in Security Education. There is a long history of using games in education (Section 4.2), and this work further attests to the benefits and value of using a game—and in this case, a physical game—in educational settings. Such games do not always match the needs of the relevant educators, but when they do match, they can provide valuable catalysts for engaging students and achieving certain learning objectives—in this case, the Awareness and Perception goals.

Game Mechanics Tradeoffs. Our main observations concern the selection of game mechanics. Overall, working with pre-existing mechanics was a positive experience, especially given our lack of expertise in the area. I wish to re-emphasize, however, the fact that mechanics directly dictate or heavily influence gameplay properties, including: how long it takes to learn to play a game; how long games take to play; the replay value of a game; and the ability to form diverse strategies. Additionally, my collaborators and I were particularly interested in how much textual content could be inserted into the game. These variables, which ultimately contribute to an

(unclearly defined) function that dictates gameplay enjoyment, are somewhat interdependent. For example, the replay value of a game is somewhat dictated by how much the game facilitates strategizing; a game's available strategies, in turn, have some relationship with the complexity of the game's rules, which directly affects the amount of time that it takes to learn a game, and partially affects the amount of time that it takes to play a game.

While these gameplay properties do not have clean-cut direct or inverse relationships, they nonetheless impact one another. When choosing or creating gameplay mechanics, sometimes tradeoffs will be necessary. It is critical to prioritize these properties in order to attempt to achieve an optimal fit.

Communication and Representation. One of the takeaways from the educator surveys was the relative importance of communicating to educators the exact nature of the game that we were distributing. While I did distribute cover letters with shipped games, they were insufficiently precise regarding the nature of the game. We never intended to design a game to teach penetration testing methods. Educators have a number of responsibilities, and may be too busy to fully vet a game before its use; it is therefore critical to provide as much information as possible regarding the nature of a game and its intended usage scenarios.

Graphic Design and Illustration. While I did not attempt to directly measure the contribution of the aesthetics of the game to achieving the goals, I do not wish to suggest its irrelevance by eliminating it from the discussion. From observation, I can comment that the graphic design, illustration, and production quality of the game seem to have a large effect at least on its initial reception. Perhaps the most poignant repeated comment that my collaborators and I have received upon presenting the game to others is, "It's like a real game!" The difference between these individuals' apparent expectations and their reaction to Control-Alt-Hack is an implicit commentary on their expectations regarding "educational games." Further study could help place the relevance of game aesthetics in the context of overall success.

4.4 Summary

My collaborators and I designed, produced, distributed, and evaluated Control-Alt-Hack: a card game designed to increase computer security awareness. The goal of this artifact is, primarily, to raise awareness of security issues and, secondarily, to improve the accuracy of people's perception of computer security as a discipline and career choice. In designing the game, we intentionally traded some technical complexity in the topics discussed in exchange for increased engagement: put another way, we set out to create a game that players could find inherently fun, from which they might learn incidentally in the course of enjoying the gameplay.

The evaluation of the game, primarily derived from the experiences of 22 educators representing over 450 students, suggests that the game accomplished its goals. Educators who used the game in their classrooms overwhelmingly indicated that they would suggest the game to others, while the majority reported both that they would use the game again and that students enjoyed the game and experienced increased security awareness. 2 educators teaching non-security computer science courses would not have taught the material without the game. A supplementary evaluation with 11 users suggested that even among a small number of participants, their reactions are aligned with a number of the goals in creating the game.

I view these results as suggesting that non-standard tools—and this game in particular—can be an effective medium for disseminating ideas and encouraging interest in computer security. In particular, given the fact that people need an increased understanding of the computer security and privacy issues surrounding emerging technologies such as implantable medical devices and augmented reality, I wanted to design this game in order to raise public awareness of new threats and risks.

4.5 Acknowledgments

This work was supported in part by Intel Labs, an Intel PhD Fellowship, the ACM Special Interest Group on Computer Science Education, and the US National Science Foundation (NSF) under awards CNS-0846065 and CNS-1247216. Any opinions, findings, and conclusions or recommendations expressed in this dissertation do not necessarily reflect the views of the funding agencies.

Control-Alt-Hack ©2012 by the University of Washington. All rights reserved. “Control-Alt-Hack” and the logo are trademarks of the University of Washington. The game mechanics are based on the game *Ninja Burger*, ©2009 by Steve Jackson Games; used under license.

Tamara Denning and Tadayoshi Kohno are founders and equity owners of RGB Hats, LLC, a private, for-profit company which has licensed the subject technology from the University of Washington. This research is subject to the conditions of a financial conflict of interest management plan established by the University of Washington.

I thank everyone who playtested, provided feedback, or otherwise helped improve the game. I thank those who graciously lent their names and biographic details: Deborah Alterman, Mike Clarke, Alexei Czeskis, Karla Danson, Iva Dermendjieva, Miro Enev, Roxana Geambasu, Sidhant Gupta, Dan Halperin, Melody Kadenko, Karl Koscher, Gabe Maganis, Cynthia Matuszek, Temitope Oluwafemi, Franz Roesner. I thank Steve Jackson for licensing his game and providing feedback, and Deborah Alterman and Mike Clarke at the University of Washingtons Center for Commercialization for helping obtain the license. I also thank and acknowledge the following people: Gravity Creative (Graphic Design), Rob Kelly (Character Illustrations), Web Design / Development (Thomas Winegarten), Juliet Fiss (Product Photos), and Tina Wegner (Production Manager).

Chapter 5

CONCLUSION

Emerging kinds of technologies are increasingly full of sensors, actuators, and wireless connectivity, and people are increasingly incorporating these technologies intimately into their lives. Security and privacy for these systems is critical; the technical and usage properties of these technologies, in combination with users' expectations, change the landscape for security and privacy.

Security and privacy is as much a human problem as a technical one; purely technically-motivated solutions all too often result in negative side effects. Human-centric methodologies must be employed to sufficiently understand an application context to design effective security and privacy.

In my thesis work, I take a human-centric approach to designing, evaluating, and promoting security and privacy for emerging technologies. Specifically, I: (1) select and refine methodologies to gather contextual information from the target application domain of implantable cardiac devices and augmented reality, resulting in recommendations for design; and (2) design, produce, distribute, and evaluate Control-Alt-Hack, a tabletop card game designed to promote awareness of high-level computer security issues.

Chapter 2 details my work in the domain of a current class of emerging technology that has physical effects in the body and is becoming increasingly interconnected: implantable cardiac devices. Implantable cardiac devices are used in a domain where their security and access control design impacts—or can be impacted by—a large number of different stakeholders including patients, nurses, cardiologists, emergency room staff, and anesthesiologists. I used mixed methods in this domain—including

inductive, qualitative methods—in order to gather information about the context of the technology usage. Results included the kinds of things that patients and medical providers value, their concerns about how problems with security might impact their jobs or their lives, and their reactions to a representative sample of different security system directions chosen from the technical security literature to embody relevant design properties. The recommendations synthesized from these studies help guide the design of security systems for future implantable cardiac devices that respect the needs, constraints, and values of multiple stakeholders.

In Chapter 3, I present my work dealing with the privacy issues surrounding a wearable emerging technology that is coming to market: augmented reality glasses. In particular, I investigate the perspectives of bystanders to augmented reality glasses regarding the technology's impact on their privacy. I use in-situ interviews to help ground participants' reactions to an emerging technology with which they may not be familiar or have pre-established opinions. Results from this work include bystanders' reactions to the glasses, the factors that they indicated make augmented reality recording different from other recording form factors, and the factors that they indicated cause augmented reality recording to be a privacy concern. Additionally, I draw inspiration from past research in this space to lay out design axes for privacy-mediating technologies and audiovisual recording.

In Chapter 4, I turn to addressing a larger issue that affects the security and privacy outcomes with emerging technologies: how to raise people's awareness of the security and privacy harms that can result from emerging technologies. In particular, I take the approach of creating a recreational, tabletop card game about computer security: Control-Alt-Hack. This chapter details the design, production, and distribution of the game, as well as an evaluation of its usage in educational contexts. Results from the evaluation suggest that the game was successful in terms of engagement and awareness and brought the security material to new audiences.

The goal of my dissertation work is to increase security and privacy outcomes

via human-centric methodologies and tools. Throughout my work, I prioritize the consideration of multiple stakeholder roles, the nuances of specific application usage, and needs and values beyond security and privacy. By incorporating these priorities—and findings that result from studies with these priorities—into system design, we can create emerging technologies that are not only secure and private, but that achieve wider acceptance, fewer negative side effects, and respect the needs and values of the people who surround them.

BIBLIOGRAPHY

- [1] [d0x3d!]. <http://www.d0x3d.com>.
- [2] DEF CON Capture the Flag. <https://www.defcon.org/html/links/dc-ctf.html>.
- [3] National Collegiate Cyber Defense Competition. <http://www.nationalccdc.org/>.
- [4] PlaidCTF. <http://play.plaidctf.com>.
- [5] Gregory D. Abowd, Gillian R. Hayes, Giovanni Iachello, Julie A. Kientz, Shwetak N. Patel, Molly M. Stevens, and Khai N. Truong. Prototypes and Paratypes: Designing Mobile and Ubiquitous Computing Applications. *IEEE Pervasive Computing*, 4(4):67–73, October 2005.
- [6] Anne Adams. Multimedia Information Changes the Whole Privacy Ballgame. In *Conference on Computers, Freedom, and Privacy*. ACM, 2000.
- [7] Erik Andersen. Optimizing Adaptivity in Educational Games. In *Proceedings of the International Conference on the Foundations of Digital Games, FDG '12*, pages 279–281, New York, NY, USA, 2012. ACM.
- [8] Sasikanth Avancha, Amit Baxi, and David Kotz. Privacy in Mobile Technology for Personal Healthcare. *ACM Comput. Surv.*, 45(1):3:1–3:54, December 2012.
- [9] Jane Bailey and Ian Kerr. Seizing Control?: The Experience Capture Experiments of Ringley & Mann. *Ethics and Inf. Technol.*, 9(2):129–139, July 2007.
- [10] Mukhtaj S. Barhm, Nidal Qwasmi, Faisal Z. Qureshi, and Khalil El-Khatib. Negotiating Privacy Preferences in Video Surveillance Systems. In *Proceedings of the 24th International Conference on Industrial Engineering and Other Applications of Applied Intelligent Systems Conference on Modern Approaches in Applied Intelligence - Volume Part II, IEA/AIE'11*, pages 511–521, Berlin, Heidelberg, 2011. Springer-Verlag.

- [11] Michael Bell and Vitali Lovich. Apparatus and methods for enforcement of policies upon a wireless device, August 28 2012. US Patent 8,254,902.
- [12] Victoria Bellotti and Abigail Sellen. Design for Privacy in Ubiquitous Computing Environments. In *ECSCW*. Kluwer Academic, 1993.
- [13] Andrew Besmer and Heather Richter Lipford. Moving Beyond Untagging: Photo Privacy in a Tagged World. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1563–1572, New York, NY, USA, 2010. ACM.
- [14] Jack Brassil. Technical Challenges in Location-Aware Video Surveillance Privacy. In Andrew Senior, editor, *Protecting Privacy in Video Surveillance*, pages 91–113. Springer London, 2009.
- [15] Bo Brinkman. The Heart of a Whistle-blower: A Corporate Decision-making Game for Computer Ethics Classes. In *Proceedings of the 40th ACM Technical Symposium on Computer Science Education*, SIGCSE '09, pages 316–320, New York, NY, USA, 2009. ACM.
- [16] Matthew Bocker and Stephen Checkoway. iSeeYou: Disabling the MacBook Webcam Indicator LED. Technical Report Technical Report 13-02, Department of Computer Science, John Hopkins University, December 2013.
- [17] William C. Cheng, Leana Golubchik, and David G. Kay. Total Recall: Are Privacy Changes Inevitable? In *Proceedings of the the 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences*, CARPE '04, pages 86–92, New York, NY, USA, 2004. ACM.
- [18] Sriram Cherukuri, Krishna K. Venkatasubramanian, and Sandeep K. S. Gupta. BioSec: A Biometric Based Approach for Securing Communication in Wireless Networks of Biosensors Implanted in the Human Body. In *ICPP Workshops*. IEEE Computer Society, 2003.
- [19] Eun Kyoung Choe, Sunny Consolvo, Jaeyeon Jung, Beverly Harrison, Shwetak N. Patel, and Julie A. Kientz. Investigating Receptiveness to Sensing and Inference in the Home Using Sensor Proxies. In *Proceedings of the 14th International Conference on Ubiquitous Computing*, UbiComp '12, pages 61–70, New York, NY, USA, 2012. ACM.
- [20] Sandy Clark, Travis Goodspeed, Perry Metzger, Zachary Wasserman, Kevin Xu, and Matt Blaze. Why (Special Agent) Johnny (Still) Can't Encrypt: A Security

- Analysis of the APCO Project 25 Two-Way Radio System. In *Proceedings of the 20th USENIX Security Symposium*, SEC'11, pages 4–4, Berkeley, CA, USA, 2011. USENIX Association.
- [21] Lorrie Cranor and Simson Garfinkel. *Security and Usability*. O'Reilly Media, Inc., 2005.
- [22] Alexei Czeskis, Ivayla Dermendjieva, Hussein Yapit, Alan Borning, Batya Friedman, Brian Gill, and Tadayoshi Kohno. Parenting from the Pocket: Value Tensions and Technical Directions for Secure and Private Parent-Teen Mobile Safety. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, SOUPS '10, pages 15:1–15:15, New York, NY, USA, 2010. ACM.
- [23] Tamara Denning, Alan Borning, Batya Friedman, Brian T. Gill, Tadayoshi Kohno, and William H. Maisel. Patients, Pacemakers, and Implantable Defibrillators: Human Values and Security for Wireless Implantable Medical Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 917–926, New York, NY, USA, 2010. ACM.
- [24] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. In Situ with Bystanders of Augmented Reality Glasses: Perspectives on Recording and Privacy-Mediating Technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2014.
- [25] Tamara Denning, Kevin Fu, and Tadayoshi Kohno. Absence Makes the Heart Grow Fonder: New Directions for Implantable Medical Device Security. In *Proceedings of the 3rd Conference on Hot Topics in Security*, HOTSEC '08, pages 5:1–5:7, Berkeley, CA, USA, 2008. USENIX Association.
- [26] Tamara Denning, Tadayoshi Kohno, and Henry M. Levy. Computer Security and the Modern Home. *Commun. ACM*, 56(1):94–103, January 2013.
- [27] Tamara Denning, Cynthia Matuszek, Karl Koscher, Joshua R. Smith, and Tadayoshi Kohno. A Spotlight on Security and Privacy Risks With Future Household Robots: Attacks and Lessons. In *Proceedings of the 11th International Conference on Ubiquitous Computing*.
- [28] Saar Drimer and Steven J. Murdoch. Keep Your Enemies Close: Distance Bounding Against Smartcard Relay Attacks. In *Proceedings of 16th USENIX Security Symposium*, SS '07, pages 7:1–7:16, Berkeley, CA, USA, 2007. USENIX Association.

- [29] Pelle Ehn and Morten Kyng. Cardboard Computers: Mocking-it-up Or Hands-on the Future. In Greenbaum and M. Kyng, editors, *Design at Work: Cooperative Design of Computer Systems*. 1992.
- [30] Ariel J. Feldman, J. Alex Halderman, and Edward W. Felten. Security Analysis of the Diebold AccuVote-TS Voting Machine. In *Proceedings of the USENIX Workshop on Accurate Electronic Voting Technology, EVT '07*, pages 2–2, Berkeley, CA, USA, 2007. USENIX Association.
- [31] Joseph L. Fleiss, Bruce Levin, and Myunghee Cho Paik. *Statistical Methods for Rates and Proportions*. John Wiley & Sons, 3rd edition, 2003.
- [32] Center for Game Science. Refraction. <http://centerforgamescience.org/portfolio/refraction/>.
- [33] Batya Friedman and Daniel C. Howe. Informed Consent in the Mozilla Browser: Implementing Value-Sensitive Design. In *Proceedings of the 35th Hawaii International Conference on System Science*, 2002.
- [34] Batya Friedman, David Hurley, Daniel C. Howe, Edward Felten, and Helen Nissenbaum. Users' Conceptions of Web Security: A Comparative Study. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems, CHI EA '02*, pages 746–747, New York, NY, USA, 2002. ACM.
- [35] Batya Friedman, Peter H. Kahn Jr., and Alan Borning. Value Sensitive Design and Information Systems: Three Case Studies. In Ping Zhang and Dennis Galletta, editors, *Human-Computer Interaction and Management Information Systems: Foundations*. 2006.
- [36] Batya Friedman, Peter H. Kahn Jr., Jennifer Hagman, Rachel Severson, and Brian Gill. The Watcher and the Watched: Social Judgements about Privacy in a Public Space. *Human-Computer Interaction*, 21(2):233–269, 2006.
- [37] Steve Jackson Games. *Ninja Burger*. 2005.
- [38] Erving Goffman. The Presentation of Self. In *Life as Theater: A Dramaturgical Sourcebook*.
- [39] Shyamnath Gollakota, Haitham Hassanieh, Benjamin Ransford, Dina Katabi, and Kevin Fu. They Can Hear Your Heartbeats: Non-Invasive Security for Implantable Medical Devices. *SIGCOMM Comput. Commun. Rev.*, 41(4):2–13, August 2011.

- [40] Mark Gondree, Zachary N. J. Peterson, and Tamara Denning. Security Through Play. *IEEE Security and Privacy*, 11(3):64–67, May 2013.
- [41] Guofei Gu, Junjie Zhang, and Wenke Lee. BotSniffer: Detecting Botnet Command and Control Channels in Network Traffic. In *Proceedings of the 15th Annual Network and Distributed System Security Symposium*, February 2008.
- [42] S. K. S. Gupta, T. Mukherjee, and K. Venkatasubramanian. Criticality Aware Access Control Model for Pervasive Applications. In *Proceedings of the Fourth Annual IEEE International Conference on Pervasive Computing and Communications*, PERCOM '06, pages 251–257, Washington, DC, USA, 2006. IEEE Computer Society.
- [43] J. Alex Halderman, Brent Waters, and Edward W. Felten. Privacy Management for Portable Recording Devices. In *Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society*, WPES '04, pages 16–24, New York, NY, USA, 2004. ACM.
- [44] Daniel Halperin, Thomas S. Heydt-Benjamin, Kevin Fu, Tadayoshi Kohno, and William H. Maisel. Security and Privacy for Implantable Medical Devices. *IEEE Pervasive Computing*, 7(1):30–39, January 2008.
- [45] Daniel Halperin, Thomas S. Heydt-Benjamin, Benjamin Ransford, Shane S. Clark, Benessa Defend, Will Morgan, Kevin Fu, Tadayoshi Kohno, and William H. Maisel. Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero-Power Defenses. In *IEEE Symposium on Security & Privacy*, 2008.
- [46] Steve Harrison and Paul Dourish. Re-place-ing Space: The Roles of Place and Space in Collaborative Systems. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work*, CSCW '96, pages 67–76, New York, NY, USA, 1996. ACM.
- [47] Benjamin Henne, Christian Szongott, and Matthew Smith. SnapMe if You Can: Privacy Threats of Other Peoples' Geo-tagged Media and What We Can Do About It. In *Proceedings of the Sixth ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec '13, pages 95–106, New York, NY, USA, 2013. ACM.
- [48] Matthew Hicks, Murph Finnicum, Samuel T. King, Milo M.K. Martin, and Jonathan M. Smith. Overcoming an Untrusted Computing Base: Detecting and

- Removing Malicious Hardware Automatically. In *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 159–172, May 2010.
- [49] Giovanni Iachello, Khai N. Truong, Gregory D. Abowd, Gillian R. Hayes, and Molly Stevens. Prototyping and Sampling Experience to Evaluate Ubiquitous Computing Privacy in the Real World. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06*, pages 1009–1018, New York, NY, USA, 2006. ACM.
- [50] Core Impact. Exploit! <http://www.coresecurity.com>.
- [51] Daniel Kahneman, Alan B. Krueger, David A. Schkade, Norbert Schwarz, and Arthur A. Stone. A Survey Method for Characterizing Daily Life Experience: The Day Reconstruction Method, 2004.
- [52] Finn Kensing and Kim Halskov Madsen. chapter Generating Visions: Future Workshops and Metaphorical Design.
- [53] Eric Klopfer, Scot Osterweil, and Katie Salen. Moving Learning Games Forward: Obstacles, Opportunities, Openness, 2009.
- [54] Karl Koscher, Alexei Czeskis, Franzi Roesner, Shwetak Patel, Tadayoshi Kohno, Stephen Checkoway, Damon McCoy, Brian Kantor, Danny Anderson, Hovav Shacham, and Stefan Savage. Experimental Analysis of a Modern Automobile. Oakland, CA, 2010.
- [55] J. Richard Landis and Gary G. Koch. The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 33:159–174, 1977.
- [56] Chunxiao Li, Anand Raghunathan, and Niraj K. Jha. Hijacking An Insulin Pump: Security Attacks and Defenses For A Diabetes Therapy System. In *e-Health Networking Applications and Services (Healthcom), 2011 13th IEEE International Conference on*, pages 150–156, June 2011.
- [57] Clara Mancini, Yvonne Rogers, Arosha K. Bandara, Tony Coe, Lukasz Jedrzejczyk, Adam N. Joinson, Blaine A. Price, Keerthi Thomas, and Bashar Nuseibeh. Contravision: Exploring Users' Reactions to Futuristic Technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10*, pages 153–162, New York, NY, USA, 2010. ACM.
- [58] Steve Mann and Joseph Ferenbok. New Media and the Power Politics of Sousveillance in a Surveillance-Dominated World. *Surveillance & Society*, 11, 2013.

- [59] Justin Manweiler, Ryan Scudellari, Zachary Cancio, and Landon P. Cox. We Saw Each Other on the Subway: Secure, Anonymous Proximity-based Missed Connections. In *Proceedings of the 10th Workshop on Mobile Computing Systems and Applications*, HotMobile '09, pages 1:1–1:6, New York, NY, USA, 2009. ACM.
- [60] M. Massimi, K.N. Truong, D. Dearman, and G.R. Hayes. Understanding Recording Technologies in Everyday Life. *Pervasive Computing, IEEE*, 9(3):64–71, July 2010.
- [61] Prabhaker Mateti. A Laboratory-based Course on Internet Security. In *Proceedings of the 34th SIGCSE Technical Symposium on Computer Science Education*, SIGCSE '03, pages 252–256, New York, NY, USA, 2003. ACM.
- [62] Microsoft. Elevation of Privilege. <http://www.microsoft.com/security/sdl/adopt/eop.aspx>.
- [63] Jessica K. Miller, Batya Friedman, and Gavin Jancke. Value Tensions in Design: The Value Sensitive Design, Development, and Appropriation of a Corporation's Groupware System. In *Proceedings of the 2007 International ACM Conference on Supporting Group Work*, GROUP '07, pages 281–290, New York, NY, USA, 2007. ACM.
- [64] Lynette I. Millett, Batya Friedman, and Edward Felten. Cookies and Web Browser Design: Toward Realizing Informed Consent Online. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '01, pages 46–52, New York, NY, USA, 2001. ACM.
- [65] David H. Nguyen, Aurora Bedford, Alexander Gerard Bretana, and Gillian R. Hayes. Situating the Concern for Information Privacy Through an Empirical Study of Responses to Video Recording. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 3207–3216, New York, NY, USA, 2011. ACM.
- [66] David H. Nguyen, Gabriela Marcu, Gillian R. Hayes, Khai N. Truong, James Scott, Marc Langheinrich, and Christof Roduner. Encountering SenseCam: Personal Recording Technologies in Everyday Life. In *Proceedings of the 11th International Conference on Ubiquitous Computing*, Ubicomp '09, pages 165–174, New York, NY, USA, 2009. ACM.
- [67] OWASP. OWASP Cornucopia Ecommerce Website Edition. https://www.owasp.org/index.php/OWASP_Cornucopia.

- [68] Leysia Palen, Marilyn Salzman, and Ed Youngs. Going Wireless: Behavior & Practice of New Mobile Phone Users. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, CSCW '00, pages 201–210, New York, NY, USA, 2000. ACM.
- [69] Shwetak N. Patel, Jay W. Summet, and Khai N. Truong. BlindSpot: Creating Capture-Resistant Spaces. In Andrew Senior, editor, *Protecting Privacy in Video Surveillance*, pages 185–201. Springer London, 2009.
- [70] Ariel Rabkin. Personal Knowledge Questions for Fallback Authentication: Security Questions in the Era of Facebook. In *Proceedings of the 4th Symposium on Usable Privacy and Security*, SOUPS '08, pages 13–23, New York, NY, USA, 2008. ACM.
- [71] Andrew Raij, Animikh Ghosh, Santosh Kumar, and Mani Srivastava. Privacy Risks Emerging From the Adoption of Innocuous Wearable Sensors In the Mobile Environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 11–20, New York, NY, USA, 2011. ACM.
- [72] Kasper Bonne Rasmussen, Claude Castelluccia, Thomas S. Heydt-Benjamin, and Srdjan Capkun. Proximity-Based Access Control for Implantable Medical Devices. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, CCS '09, pages 410–419, New York, NY, USA, 2009. ACM.
- [73] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. Detecting and Defending Against Third-party Tracking on the Web. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*, NSDI'12, pages 12–12, Berkeley, CA, USA, 2012. USENIX Association.
- [74] Michael Rushanan, Colleen Swanson, Denis Foo Kune, and Aviel D. Rubin. SoK: Security and Privacy in Implantable Medical Devices and Body Area Networks. In *IEEE Symposium on Security & Privacy*, 2014.
- [75] Stuart E. Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. The Emperor's New Security Indicators. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, SP '07, pages 51–65, Washington, DC, USA, 2007. IEEE Computer Society.
- [76] Stuart Schechter. Security That Is Meant to Be Skin Deep: Using Ultraviolet Micropigmentation to Store Emergency-Access Keys for Implantable Medical Devices. In *USENIX Workshop on Health Security and Privacy*, HealthSec '10. USENIX, 2010.

- [77] Jeremy Schiff, Marci Meingast, Deirdre K. Mulligan, Shankar Sastry, and Ken Goldberg. In *International Conference on Intelligent Robots and Systems*. IEEE/RSJ.
- [78] Katie Seaborn, Magy Seif El-Nasr, David Milam, and Darren Yung. Programming, PWNed: Using Digital Game Development to Enhance Learners' Competency and Self-efficacy in a High School Computing Science Course. In *Proceedings of the 43rd ACM Technical Symposium on Computer Science Education, SIGCSE '12*, pages 93–98, New York, NY, USA, 2012. ACM.
- [79] Steve Sheng, Bryant Magnien, Ponnurangam Kumaraguru, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *Proceedings of the 3rd Symposium on Usable Privacy and Security, SOUPS '07*, pages 88–99, New York, NY, USA, 2007. ACM.
- [80] Boris Shimanovsky, Jessica Feng, and Miodrag Potkonjak. Hiding Data in DNA. In *Revised Papers from the 5th International Workshop on Information Hiding, IH '02*, pages 373–386, London, UK, UK, 2003. Springer-Verlag.
- [81] Adam Shostack. Elevation of Privilege: Drawing Developers into Threat Modeling. Technical report, Microsoft, 2012.
- [82] H. Jeff Smith and Sandra J. Milberg. Information Privacy: Measuring Individuals' Concerns About Organizational Practices. *MIS Q.*, 20(2):167–196, June 1996.
- [83] Jacob Sorber, Minh Shin, Ronald Peterson, Cory Cornelius, Shrirang Mare, Aarathi Prasad, Zachary Marois, Emma Smithayer, and David Kotz. An Amulet for Trustworthy Wearable mHealth. In *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications, HotMobile '12*, pages 7:1–7:6, New York, NY, USA, 2012. ACM.
- [84] Anna C. Squicciarini, Heng Xu, and Xiaolong (Luke) Zhang. CoPE: Enabling Collaborative Privacy Management In Online Social Networks. *Journal of the American Society for Information Science and Technology*, 62(3):521–534, 2011.
- [85] Thad Starner, Steve Mann, Bradley Rhodes, Jeffrey Levine, Jennifer Healey, Dana Kirsch, Rosalind W. Picard, and Alex Pentland. *Augmented Reality Through Wearable Computing*, 1997.

- [86] Latanya Sweeney. Weaving Technology and Policy Together to Maintain Confidentiality. *Journal of Law, Medicine & Ethics*, 25, 1997.
- [87] The Center for Information Systems Security Studies and Research, Naval Postgraduate School. CyberCIEGE. <http://cisr.nps.edu/cyberciege/>.
- [88] US Department of Defense. CyberProtect. <http://iase.disa.mil/eta/cyber-protect/launchpage.htm>.
- [89] Krishna K. Venkatasubramanian, Ayan Banerjee, and Sandeep Kumar S. Gupta. PSKA: Usable and Secure Key Agreement Scheme for Body Area Networks. *Trans. Info. Tech. Biomed.*, 14(1):60–68, January 2010.
- [90] Rick Wash. Folk Models of Home Computer Security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, SOUPS '10, pages 11:1–11:16, New York, NY, USA, 2010. ACM.
- [91] WeWantToKnow. DragonBox. <http://www.dragonboxapp.com/>.
- [92] Georgory White and Georgory Nordstrom. In Dorothy E. Denning and Peter J. Denning, editors, *Internet Besieged*, chapter Security Across the Curriculum: Using Computer Security to Teach Computer Science Principles, pages 519–525. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 1998.
- [93] Alma Whitten and J. D. Tygar. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. In *Proceedings of the 8th conference on USENIX Security Symposium*, SSYM'99, pages 14–14, Berkeley, CA, USA, 1999. USENIX Association.
- [94] Laurie Williams, Andrew Meneely, and Grant Shipley. Protection Poker: The New Software Security “Game”. *IEEE Security & Privacy*, 8(3):14–20, 2010.
- [95] Eric Wustrow, Scott Wolchok, Ian Goldberg, and J. Alex Halderman. Telex: Anticensorship in the Network Infrastructure. In *Proceedings of the 20th USENIX Conference on Security*, SEC'11, pages 30–30, Berkeley, CA, USA, 2011. USENIX Association.
- [96] Fengyuan Xu, Zhengrui Qin, C.C. Tan, Baosheng Wang, and Qun Li. IMDGuard: Securing Implantable Medical Devices with the External Wearable Guardian. In *INFOCOM, 2011 Proceedings IEEE*, pages 1862–1870, April 2011.

- [97] Daisy Yoo, Milli Lake, Trond Nilsen, Molly E. Utter, Robert Alsdorf, Theoneste Bizimana, Lisa P. Nathan, Mark Ring, Elizabeth J. Utter, Robert F. Utter, and Batya Friedman. Envisioning Across Generations: A Multi-lifespan Information System for International Justice in Rwanda. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13*, pages 2527–2536, New York, NY, USA, 2013. ACM.
- [98] Thomas Zimmerman. Personal Area Networks: Near-Field Intrabody Communication. *IBM Systems Journal*, 35(3 & 4), 1996.

Appendix A

EXTENDED CATEGORY DETAILS FOR MEDICAL PROVIDER WORKSHOP STUDY

A.1 Initial Perspectives

Section 2.4.2 discusses results from the data analysis of the Initial Perspectives worksheets completed by participants during the workshop. Table A.1 presents: definitions for the categories that emerged during data analysis, example responses, and the percentage of participants who expressed the issue on their worksheet.

Initial Perspectives Categories	Definitions and Examples	Participant Percentages N = 24
Access & Sharing	Refers to the importance of the provider's access to (and the sharing of) data, devices, and individuals with expertise. (E.g.; <i>"the ability to access device data and easily transmit that information to the electronic medical record or share the info with experts in a remote location"</i> ; <i>"ability to interrogate prior events & analyze rhythms detected"</i>)	25
Compatibility	Refers to the importance of inter-manufacturer, inter-provider, and inter-device compatibility with respect to identifying implantable devices, programming them, terminology, and settings. (E.g., <i>"idiosyncrasies of device setting restoration among companies + even devices"</i> ; <i>"the ability to interrogate devices with a universal interrogation device..."</i>)	29
Correct Usage	Refers to the importance of correctly identifying which individuals should have an implantable cardiac device, and a medical provider having the expertise to choose appropriate settings for the device given the patient's condition. (E.g., <i>"iatrogenic harm to patient from inappropriate programming or programming choices"</i> ; <i>"value of life saving therapy in appropriate patient"</i>)	25
Device Battery Life	Refers to the importance of maximizing the battery lifetime of a device, or developing rechargeable solutions. (E.g., <i>"battery longevity or rechargeability"</i> ; <i>"battery life"</i>)	17

Device Compactness / Inertness	Refers to the importance of the implantable medical device having a small form factor and minimal effects on other medical diagnostic or therapeutic procedures. (E.g., “ <i>can have issues w/ getting MRI or imaging needed b/c of devices</i> ”; “ <i>NO bigger → smaller the better</i> ”)	21
Device Ecosystem	Refers to the importance of (non-compatibility-related) issues in the larger device ecosystem such as investment, innovation, regulation, or testing. (E.g., “ <i>this is one of the most innovative investment space [sic]</i> ”; “ <i>regulatory oversight</i> ”)	13
Device Functionality	Refers to the importance of the device’s malfunction-free operation for the purpose for which it was implanted, including the device’s therapies and related functions, general device and lead reliability, and low false positive and false negative rates for ICD shocks. (E.g., “ <i>malfunction (inappropriate shocks)</i> ”; “ <i>they work to accomplish the medical task they were implanted for</i> ”)	79
Patient / Patient Health	Refers to the importance of the patient with regards to health, mental wellbeing, quality of life, and comfort with and understanding of the medical treatment. (E.g., “ <i>they are lifesaving devices for patients with dysrhythmias</i> ”; “ <i>[patients] who have had shocks complain of fear of future shocks</i> ”)	75
Programming	Refers to the importance of an easy, understandable, and productive experience with the programmer (equipment). (E.g., “ <i>ease of interface (programmer)</i> ”; “ <i>complete physician programmability, i.e. avoid automatic device decision making</i> ”)	25
Quality of Data	Refers to the importance of accurate and precise data. (E.g., “ <i>detailed, accurate info re: heart rhythm for diagnostics</i> ”)	4
Remote Monitoring	Refers to the importance of the ability to remotely monitor a patient’s status via the remote monitoring system. (E.g., “ <i>ability to be remotely monitored</i> ”; “ <i>remote monitoring (of rhythm, volume stats)</i> ”)	21
Security & Privacy	Refers to the importance of security and privacy with respect to personal health information and medical devices. (E.g., “ <i>security, privacy, confidentiality of personal informatics</i> ”; “ <i>some modest level of security is lacking</i> ”)	8
Surgery & Healing	Refers to the importance of an easy and successful surgical process with infection-free recovery. (E.g., “ <i>life-threatening complication from implantation</i> ”; “ <i>complications of ICD including infection</i> ”)	58

Uncodable	Refers to a response that is uncodable due to ambiguity or illegibility. (E.g., “ <i>safety</i> ”; “ <i>human factors related to implants</i> ”)	25
-----------	--	----

Table A.1: Percentage of participants who mentioned an issue on their Initial Perspectives worksheet, by emergent analysis category. Shaded cells represent categories used by more than 50% of the participants.

A.2 Metaphors

Section 2.4.2 discusses results from the data analysis of the metaphors contributed by participants during the workshop. Table A.2 presents: definitions for the categories that emerged during data analysis, example responses, and the percentage of participants who expressed the issue on their worksheet.

Metaphor Categories	Definitions and Examples	Metaphors	
		Implantable Cardiac Device [IMD] Total = 42	Security [SEC] Total = 39
Agency	Refers to an agent or qualities of agency (e.g., “ <i>guardian angel</i> ” [IMD]; “ <i>personal paramedic</i> ” [IMD]).	3	0
Bio-medical	Refers to biological processes, medical practices, and the preservation of life (e.g., “ <i>takes up half the chest wall</i> ” [IMD]; “ <i>impact on other medical care</i> ” [SEC]).	10	1
Business	Refers to motivations, practices or bureaucracy of large organizations (e.g. “ <i>profit for device company or hospital</i> ” [IMD]; “ <i>proprietary closed systems</i> ” [SEC]).	2	2
Emotion	Refers to affective responses (e.g., “ <i>anxiety provoking</i> ” [IMD]; “ <i>whose fear?</i> ” [SEC]).	6	2
Information	Refers to control and flow of information (e.g., “ <i>hub of information</i> ” [IMD]; “ <i>controlling the information</i> ” [SEC]).	2	4
Maintenance	Refers to required upkeep and on-going medical follow up (e.g., “ <i>lifelong issue for patient</i> ” [IMD]; “ <i>complicating to the management of devices</i> ” [SEC]).	4	1

Personal Identity	Refers to a person's sense of self and wholeness as a person (e.g., " <i>tin mans heart</i> " [IMD]; " <i>tangibly belongs to one person</i> " [SEC]).	3	1
Privacy	Refers to personal information and information privacy (e.g., " <i>teenager with inappropriate Facebook photos exposed to the world</i> " [SEC]).	0	3
Risk	Refers to assessing, managing, and balancing risk factors (e.g., " <i>safety net</i> " [IMD]; " <i>how real is the risk?</i> " [SEC]).	3	7
Security	Refers to presence or absence of security and security systems (e.g., " <i>anti-virus protection</i> " [SEC]; " <i>bank with an unlocked vault</i> " [SEC]).	0	14
Technology	Refers to mechanical or electronic artifacts (e.g., " <i>engine</i> " [IMD]; " <i>wireless router</i> " [SEC]).	3	2
Uncodable	Refers to ambiguous items (e.g., " <i>monitor</i> " [IMD] could be monitoring the technology or monitoring biological systems) or items that do not fit into one of the above categories (e.g., " <i>wave of the future</i> ").	6	2

Table A.2: Numbers of group-generated metaphors, by emergent analysis category. Given are numbers of metaphors generated during the sections on metaphors of implantable cardiac devices (IMD) and metaphors of security for implantable cardiac devices (SEC), respectively. The categories are alphabetical.

Appendix B

ADDITIONAL CONTROL-ALT-HACK MATERIAL

B.1 Card Topics and Research Papers

Table B.1 gives some examples of Mission cards that were inspired by research results. These examples list one relevant research project per Mission; I acknowledge that other examples exist and that this is not a comprehensive list.

Card Title	Card Topic	Example Inspirational Research
[CENSORED]	Working on steganographic anti-censorship software	[95]
A Healthy Dose of Security	Consulting to improve the security of an insulin pump	[56]
A Rash Decision	Cross-correlating data sources to de-anonymize medical records	[86]
Cookie-Blocked	Writing a web browser extension to circumvent tracking cookies	[73]
Crash Test Dummy	Hacking an automobile	[54]
<i>E. coli</i> Cryptography	Implementing cryptography via synthetic biology	[80]
Hay Baby, Hay Baby, Hay	Demonstrating that a dating site has insecure password recovery questions	[70]
Here's Looking at You, Kid	Analyzing the security of a WiFi-enabled, webcam-equipped toy robot	[27]
I'd Tap That	Pen testing the security of a contactless payment system	[28]
Mr. Botneto	Measuring a botnet's growth, then reverse engineering the C&C algorithm	[41]
One Hacker, Won Vote	Pen testing an electronic voting machine	[30]
Trojan Protection	Looking for backdoors in the outsourced production of hardware	[48]

Table B.1: Example Mission card titles, topics, and example research that inspired them.

B.2 Educator Survey Contents

The questions asked in the online survey distributed to educators are given below:

1. How did you use Control-Alt-Hack®? Please describe the activity.
2. How long did the activity take?
3. Were there any written or oral components that students turned in or presented as part of the activity? If so, please describe.
4. Did you present or assign any supplementary materials? If so, please describe.
5. What, if anything, worked well with the activity?
6. What, if anything, would you do differently if you were to do the activity again?
7. How would you describe students' level of enjoyment and/or engagement with the activity?
8. How would you describe students' level of learning with the activity? On what particular topics was their learning focused?
9. Why did you choose to use Control-Alt-Hack® in your classroom?
10. Would you use Control-Alt-Hack® again in your classroom?
11. Why or why not?
12. Would you suggest Control-Alt-Hack® to others?
13. Why or why not?

14. If you had not used Control-Alt-Hack® , would you still have covered the material?
15. Did you cover the material using another (additional) method?
16. If applicable: What additional method did you use to cover the material?
17. If applicable: How would you compare these two methods (Control-Alt-Hack® and the additional method) of covering the material? What are the pros of each? What are the cons?
18. If applicable: If you had not used Control-Alt-Hack® , what alternative method would you have used to cover the material?
19. If applicable: How would you compare these two methods (Control-Alt-Hack® and the alternative method) of covering the material? What are the pros of each? What are the cons?
20. What is the subject of your class?
21. What is the class format (e.g., MWF 50-min 10-week course, 2-hour training seminar, etc.)?
22. How many students participated in the activity?
23. What is the level of the students in your class?
24. What is the (approximate) level of student experience with computer science and/or computer security?
25. Is there anything else that you would like to add that we have not addressed?

B.3 Educator Survey Coding Disagreements

I include all 11 cases where the primary coder's and the secondary coder's coding results did not agree. The examples are given below, along with the quotes from the survey which were primarily responsible for the distinction between the coding results.

As mentioned in Section 4.3.1, except for one case in which the secondary coder misread the data and coded an error (Case 6), the primary coder's results—the results reported in the paper—always represent the stricter of the two viewpoints. That is, in Section 4.3.2 I report the upper bound on the coders' interpretation of the critiques to the game and the lower bound on the game's role in engagement and awareness.

B.3.1 Positive Role

Below I provide information on the cases where the primary and reliability coder disagreed when coding the positive role(s) that the game performed. In all cases, the reliability coder coded the game as playing the role, while the primary coder did not. Quotes that led the reliability coder to code the game as “Social / Engagement” or “Awareness” are given below.

Social / Engagement:

- Case 1 (E11-checkout). *“It was a 7/10. the students enjoyed it but the word did not spread around and ignite students.”*
- Case 2 (E12-classroom). *“Students reported that they enjoyed the game, but that the hour twenty was pushing the limit.”*
- Case 3 (E16-classroom). *“The kids were all engaged with the game and playing it through.” “Would rate it 8/10.”*

- Case 4 (E22-vetting). *“They seemed engaged, although not so much that I would expect they would play it for fun.”*

Awareness:

- Case 5 (E21-lunch). *“Brought up some terminology that staff and IT had not heard before. “pwned” :-)”*

B.3.2 Critiques and Tradeoffs

Below we provide information on the cases where the primary and reliability coder disagreed when coding critiques made to the game. In all cases except one (Case 6, coded in error), the primary coder coded the educator as offering that critique, while the reliability coder did not. Quotes that led primary coder to code the critique are given below.

Takes a long time to learn:

- Case 6 (E10-classroom): The disagreement was due to the reliability coder misreading the response. The reference to the presentation and the gameplay taking too long together was a reference to instructor’s syllabus content, not the video introducing the game’s rules.

Takes a long time to play:

- Case 7 (E1-classroom): [Q: What, if anything, would you do differently if you were to do the activity again?] *“have more play time during the topic”*

Not enough fun:

- Case 8 (E22-vetting): *“They seemed engaged, though not so much that I would expect them to play it for fun”*

Not enough educational value:

- Case 9 (E2-ACM): *“Learning was not so much learned throughout the game, but it did pose interesting questions that the students were curious about”*
- Case 10 (E3-vetting): *“I worry the card game will seem like a card game”*
- Case 11 (E12-classroom): *“Most students reported a low level of learning, the topics that were reported positively were presenting the students with real world context for what they were learning”*

B.4 Control-Alt-Hack-themed Assignments

5 of the 14 educators who used Control-Alt-Hack in the classroom reported using a custom assignment in concert with the game, as described below:

E4-classroom (12 undergraduate students with little or no prior security experience): Students were asked to identify at least three tasks from Mission cards that seemed interesting. A follow-up exercise may be to have them research real-life situations where the theme of one of the tasks is involved.

E7-classroom (56 undergraduate students in a Cyber-Security and Information Assurance course, with little or no prior security experience): Students were asked to take a scenario from a card and craft a research paper inspired by the scenario.

E10-classroom (65 high school students in a Game Design course with little or no security background): Students were required to answer essay questions about the game and how it is put together. Optional questions asked about the game’s relation to the IT industry and hacker culture.

E12-classroom (22 undergraduates in a Computer Security Course with prior educational security experience): Students wrote one to two paragraphs discussing the activity.

E18-classroom (30 undergraduate students in a Fundamentals of Information Security course with a variety of security backgrounds). Two

questions were asked before the game: (1) *What does information security mean to you?*; and (2) *What skills are required in white-hat hacking?* Two questions were asked after the game: (1) *Did your answers to the previous questions change as a result of the game—and if so, how?*; and (2) *As a result of the game, did you discover any threats you hadn't considered—and if so, what?*

While some of the above assignments are similar to activities we propose on our web site (<http://www.controlalthack.com>), some of the assignments are original and demonstrate an interesting integration of the game into existing course plans and practices.