Privacy in Mobile Technology for Personal Healthcare

SASIKANTH AVANCHA and AMIT BAXI, Intel Labs Bangalore
DAVID KOTZ, Dartmouth College

Information technology can improve the quality, efficiency, and cost of healthcare. In this survey, we examine the privacy requirements of mobile computing technologies that have the potential to transform healthcare. Such *mHealth* technology enables physicians to remotely monitor patients' health and enables individuals to manage their own health more easily. Despite these advantages, privacy is essential for any personal monitoring technology. Through an extensive survey of the literature, we develop a conceptual privacy framework for mHealth, itemize the privacy properties needed in mHealth systems, and discuss the technologies that could support privacy-sensitive mHealth systems. We end with a list of open research questions.

1. INTRODUCTION

Healthcare information technology (IT) has huge potential to improve healthcare quality, improve efficiency, and reduce cost, and is currently on the cusp of major innovations and widespread deployment in the US and elsewhere. Its potential may best be described by a quote from the chairs of three leading healthcare policy and standards groups in the US.

> "Our vision is one of a 21st century health system in which all health information is electronic, delivered instantly and securely to individuals and their care providers when needed, and capable of analysis for constant improvement and research. With better information upon which to base decisions, the challenging process of health reform can successfully proceed—measuring

quality, rewarding value, engaging individuals—and lead the way to better health for all Americans." [Halamka et al. 2009]

In this survey, we specifically examine the privacy challenges involved in *mobile* computing and communications technologies. Such mHealth technology [mH 2009] appears promising in many ways: enabling physicians to remotely monitor their patients' health and improve the quality of healthcare, enabling patients to manage their health more easily, and reducing the cost of care by allowing patients to spend less time in the hospital or make fewer visits to their doctor.

In mHealth, Mobile Internet Devices (MIDs), connected wirelessly to wearable, portable, and even embeddable sensors, will enable long-term continuous medical monitoring for many purposes [Baker et al. 2007; Boric-Lubecke and Lubecke 2002; Varshney 2007]: for outpatients with chronic medical conditions (such as diabetes), individuals seeking to change behavior (such as losing weight), physicians needing to quantify and detect behavioral aberrations for early diagnosis (such as depression), or athletes wishing to monitor their condition and performance. In this article, we use the term "Patient" to describe the subject of sensing in all such use cases, using the capitalized form as a reminder of its broader meaning. We expect MIDs, such as smart phones, to contain the technology and applications needed to process sensor data and enable their appropriate use. The resulting data may be used directly by the Patient [AC 2008; AH 2008; Wang et al. 2006] or may be shared with others: with a physician for treatment [SH 2008], with an insurance company for coverage, with a scientist for research [DH 2008], with a coach for athletic training [Aylward and Paradiso 2007], or with family members and friends in social-networking communities targeted towards health and wellness [OW 2009; DS 2009, e.g.].

The term "mHealth" applies broadly to the use of mobile technology in healthcare applications. In this article, however, we focus on patient-centered technology, as described in the previous examples and the detailed scenarios that follow. There are, of course, valuable uses of mobile technology in other aspects of healthcare delivery and management, including personal communication devices used by clinicians, inventory-control systems for medical equipment and consumables, and telemedicine platforms for emergency response or remote rural healthcare. Some of the issues we raise occur in such settings, but we do not directly address them here.

1.1. The Challenge

Although mHealth systems may indeed improve quality of healthcare and quality of life, they also generate new security and privacy issues [Al Ameen et al. 2010; Giannetsos et al. 2011]. The technology goal should be to develop usable devices that respect Patient privacy while also retaining the data quality and accessibility required for the medical uses of the data. In this article, we focus on privacy; specifically, we wish to give the Patient control over the data that is collected and to whom it is disclosed, and to recognize that different situations may require different responses. Indeed, we note that control, not possession or ownership, is fundamental to privacy.[1] Privacy means that the Patient retains control even when the data is "owned" by another party (as is common in medical records maintained by a hospital) and even after a copy of the data has been provided to another party (as when billing records are shared with an insurance company).

Given our focus on privacy, it is essential that we define it clearly for the context of healthcare. Fortunately, others have thought deeply about this issue; we adopt the

---

[1]Some researchers in the field of healthcare information privacy have a different opinion [Nissenbaum 2004], but the notion of privacy as information-control is common and (as we show in the next paragraph) the core of the definition used by an important advisory committee.

definition selected by the National Committee for Vital and Health Statistics (NCVHS), a key advisory committee to the US Department of Health and Human Services. "*Health information privacy* is an individual's right to control the acquisition, uses, or disclosures of his or her identifiable health data. *Confidentiality*, which is closely related, refers to the obligations of those who receive information to respect the privacy interests of those to whom the data relate. *Security* is altogether different. It refers to physical, technological, or administrative safeguards or tools used to protect identifiable health data from unwarranted access or disclosure" [Cohn 2006]. We also follow NCVHS and define PHI as "personal health information" rather than "protected health information", which is a phrase that has specific meaning in a HIPAA context [Choi et al. 2006].

Clearly, privacy is important in any healthcare information system. What is different or especially challenging about mHealth privacy? First, mHealth allows for the collection of far more medical data about the Patient, as many mHealth devices collect data continuously over extended periods of time. (For example, it is possible to record ECG data continuously for weeks, throughout daily life, rather than a one-minute recording taken in the clinic every other week.) Second, mHealth allows much broader range of health-related information to be collected, not just physiological data; many mHealth applications will collect information about Patient lifestyle and activities (such as food habits and diet details, location tracks, physical activity, or social interactions). Third, mHealth will enable a broad range of health-related applications: sharing data with your health provider, as in a traditional doctor relationship, but also sharing data with an insurance company (e.g., to confirm compliance with a medication regimen), with lifestyle coaches (e.g., diet advisers), with athletic coaches (e.g., sports teams or health-club trainers), or with family (e.g., to support a relative's recovery from surgery). In such settings, privacy is a complex issue: the Patient needs subtle control over the collection, recording, dissemination, and access to their mHealth data.

1.2. Background

In many cases, the data collected by mHealth devices will be incorporated into a medical record. Information from a medical record is sought by and may be shared with several players: physicians, caregivers, medical experts, family, hospitals, insurers, clearing-houses, laboratories, pharmacies, government agencies, employers, managed care organizations, academic research organizations and public health organizations.

There are at least two broad categories of medical records. An Electronic Health Record (EHR) is created and managed by a healthcare provider (hospitals and other clinical organizations), whereas a Personal Health Record (PHR) is created and managed by the Patient herself. Since PHRs pose at least as many privacy challenges as EHRs, we focus primarily on PHRs in this article.

Google Health [GH 2008] and Microsoft's HealthVault [MHV 2008] are two well-known PHR services. These systems allow people to easily access their PHI via a web portal, allowing Patients to collate PHI from various sources, and perform other operations such as deleting, editing, and sharing their PHI with multiple entities including family, friends, and health care professionals. People can manage PHRs of family members in one Google Health or Microsoft HealthVault account, enabling co-management of PHRs by designated family members or friends. These PHR services have relationships with other health service providers that offer services to the user. Users are allowed to share information in their PHRs at any level of granularity they choose, with external health service providers, caregivers, coaches, trainers and doctors.

Several custodial PHR models are evolving [CHCF 2008], which vary in who manages and controls the health information. A healthcare provider manages PHRs for hospitals in the *provider-based PHR* model; Patients can view their records, schedule

appointments, and consult physicians via web portals. In this model, the hospital has greater contrl over PHI than the Patient. In the *health-plan* or *employer-based PHR* model, Patients may have greater control over PHI and can also obtain information related to claims and lab results from multiple providers, however, clinicians continue to exercise control over PHI in their possession. In the *patient-centric PHR* model, Patients control all their PHI via web portals or portable media devices, including operations to create, import, update, read and delete records. This model increasingly allows compatible devices to upload PHI directly to Patient PHR via appropriate interfaces. Another common case is the *vendor-supplied PHR*, in which an mHealth-device vendor provides an application-specific record of the data collected by that device, accessible to the Patient on the vendor's website. The mHealth-related privacy issues in such a system are the same as those in the patient-centric PHRs mentioned previously. We expect Patients will be challenged, however, to manage their privacy across multiple PHRs, and to understand the subtle complexities of their trust relationships with vendors, wireless carriers, Internet service providers, and healthcare providers.

As the use of patient-centric PHR services via mobile sensing, aggregation and communication devices increases, PHR service providers will be faced with new problems related to ensuring user privacy. Several problems arise due to vulnerabilities associated with mobile sensors and mobile nodes; we delve into these problems in the following text. Fortunately, mobile-phone platforms are becoming more open, reducing the control of cellular carriers (and the privacy risks inherent in trusting them with health data); on the other hand, open mobile-phone platforms may be more vulnerable to Internet-based threats.

1.2.1. Architecture and Terminology. Because this topic includes a lot of terminology and many acronyms, we provide a quick-reference glossary in Appendix B.

We imagine an infrastructure in which each Patient carries a mobile node (MN), which may be their mobile phone or other mobile Internet device (MID), and a personal collection of sensor nodes (SNs) that can measure data about their activity (accelerometers, pedometers, location) or physiology (electrocardiograms, pulse oximeters, blood-glucose meters, weight scales). These sensors may be carried by the Patient [Mack et al. 2006], worn by the Patient [Paradiso et al. 2005], embedded in their living space [SH 2008], or implanted in their body [Halperin et al. 2008b]. The sensors communicate with the MN through a wireless body-area network. The MN is responsible for coordinating the sensors, collecting the sensor data, (optionally) aggregating or pre-processing the sensor data, and reporting the data to a health records system (HRS). The MN also serves as the Patient's primary interface to the HRS, with respect to managing the data-collection process and subsequent sharing of the data.

The health records system (HRS) may be an Electronic Health Record (EHR) or Personal Health Record (PHR) system. Most of the mobile-centric research challenges will be the same for both EHR and PHR scenarios. We must also consider the likely situation, over the next decade, in which a Patient must interact with multiple health-records systems from multiple providers and built on different models.

The Consumers of these records (including doctors and other clinical personnel, insurance companies and other billing-related personnel, researchers and regulators) access the HRS through some Client computer. The security issues on this platform are largely out of scope of this article, except in cases where we seek end-to-end technical mechanisms to support the Patient's privacy wishes.

Finally, these people and systems need a supportive ecosystem, a set of authorities and agencies that provide the regulatory, logistical, and technical foundation for the above relationships. We can identify at least five roles to be played by some combination of public and private organizations.

—*Policymakers* establish laws, regulations, and standards regarding the protection of Patient privacy in mHealth technology; ultimately, federal and state agencies enforce these laws and regulations.

—*Certification bodies* attest to whether particular products and services meet the policies and standards.

—*Manufacturers* produce hardware and software products and services, such as the MNs, SNs, and HRS.

—*Distribution & Management* services distribute the hardware and software to Patients and Consumers, and provide remote-management capabilities such as secure, automatic software updates and remote deletion of data and keys on lost devices.

—*Key-Management Infrastructure* provides the key-distribution and certificate authorities to support the crypto-systems used for verification (e.g., to verify the signature of a certification body regarding an SN calibration, or to verify the public key of a management service).

We should take care not to expect that a top-down definition of, let alone deployment of, such an infrastructure is possible. Any such system will necessarily be defined and deployed by many organizations and agencies, over time [Rouse 2008].

In this article, we focus on the *mobile* aspects of the infrastructure, in which there are many risks: the sensor data may be intercepted (impacting privacy), tampered with (leading to incorrect data and care decisions), or blocked (leading to loss of information to researchers or care providers). Furthermore, MNs or SNs may be lost or stolen, resulting in possible exposure of any data or encryption keys they contain. Finally, since we expect that Patients would like to use their existing mobile phone as their MN, these risks are compounded because health-sensing tasks must share the phone with email, web browsing, and other activities that open the platform to risk of compromise. All of these risks can lead to dangerous consequences for the Patient and provider [Kulkarni and Öztürk 2007; Stanford 2002].

Although any viable solution, and any real deployment, will doubtless be more complex than implied by the above description, this architecture provides a structural basis and terminology for our discussion of prior work and upcoming research challenges, below.

1.3. Contributions

We make four broad contributions in this article.

(1) We describe a conceptual privacy framework for the mHealth context.
(2) We articulate the properties required of a privacy-sensitive mHealth solution.
(3) We survey prior work, identifying what ideas and technologies could be adopted.
(4) We identify the key research challenges that the community should address.

We begin with a brief summary of mHealth scenarios, which help to set the context. In the next section, we survey the legal and ethical foundations of privacy as they apply to healthcare information technology, examining several conceptual privacy frameworks. In Section 3 we describe our conceptual privacy framework for mHealth, and itemize the necessary privacy and security properties. These properties should be useful to anyone designing such systems. We then take a moment to explore a case study in Section 4, before returning in Section 5 to look at the technology relevant to privacy in healthcare and in mobile systems. Section 6 discusses challenges for the research community, representing gaps between the desired properties and the available technology. After the summary in Section 7, we include two appendices: a detailed listing of healthcare privacy frameworks in Appendix A, and a glossary in Appendix B.

1.4. mHealth Scenarios

To better set the context for mHealth applications and their privacy challenges, we describe a set of scenarios ranging from emergency situations to illness management.

Scenario 1—Emergency Health Situation. Anna suffers from cardiovascular disease, and has infrequent episodes of abnormal heart rhythms (arrhythmias) that may be dangerous. Because arrhythmia can occur anytime, anywhere, Anna needs constant monitoring to detect the condition when it occurs. To detect an arrhythmia, wearable sensors monitor ECG patterns and heart rate and an alarm is raised if there is a significant deviation from normal state. However, ECG pattern and heart rate may also change due to physical activity. To distinguish between an arrhythmia and a normal increase in heart rate induced by physical activity, Anna's mobile phone includes a GPS and a physical activity monitor, and interprets the ECG data in the context of those sensors' data. During one of her morning jogs in the park, Anna's mobile device begins to beep. Anna looks at the screen, which indicates the onset of an arrhythmia. Even as Anna begins to sit down, she starts feeling the effects of the attack. Because Anna did not turn off the alarm within a specific period, the device sends an alert to Anna's health-monitoring service via the cellular network. The dispatcher sees the alarm on his patient-tracking console, dispatches the nearest available mobile caregiver to Anna's location, and notifies her primary-care provider. In this scenario, Anna has two privacy concerns. First, Anna wants the mobile device to share her location only with the monitoring service and only in an emergency. Second, she would not like others nearby to be aware of her heart condition. (Scenario adapted from Jones et al. [2007].)

Scenario 2—Post-Surgical Home Monitoring. Following minor surgery, Jane returns home after recuperating in the hospital for one day. Because Jane's healthcare management organization offered to reduce her premiums if she opted for in-home monitoring instead of in-hospital monitoring following minor surgeries, Jane leases a suite of sensors, including ECG, blood oxygenation, and blood pressure. Jane applies these sensors to her body during her day's stay in the hospital. With the help of the nurse, Jane ensures that her smart mobile phone is capturing her vitals from these sensors periodically and transmitting the data to her electronic medical record (EMR) at the hospital.

At Jane's discharge meeting, the surgeon instructs Jane's husband Jim in the proper application of the sensors and loads an app on Jim's smart-phone so that he can view data collected by the sensors, and record observations (such as photographs of the wound healing) and events (such as medications given). With a few taps on Jim's phone, Jane grants her consent for Jim's phone to link to her sensors and to her EMR, allowing periodic upload of the home-collected data; later, the surgeon can view the EMR data if Jane or Jim calls with any questions. This consent step supports Jane's privacy by giving her control over the collection and sharing of the sensor data. (Scenario adapted from Sriram et al. [2009b].)

Scenario 3—Diabetes Management. Ravi is a diabetic who finds it difficult to manage his condition effectively, resulting in significant variation of his diurnal blood-glucose levels. Ravi's doctor advises him to subscribe to a Diabetes Management Program offered by his hospital. As a part of the program, Ravi wears a hospital-provided device that continuously monitors his activity level and calories burned, and installs software on his mobile phone. The software processes data it receives from the monitor along with contextual information such as Ravi's location, calendar schedule and time of the day. It reminds him to take his medication, alerts him to long periods of inactivity, encourages him to log diet information and tracks his daily goals of calorie intake and expenditure. The mobile phone periodically synchronizes with Ravi's PHR.

Ravi decides to join a social network for diabetics, whose privacy settings allow him to share information with the group, such as his daily activity, food intake progress and goals. Ravi also has the choice to remain anonymous in the group, in which case he would only know his ranking (based on compliance with activity, diet and medications) relative to others in group. Separately from the social network, Ravi chooses to allow complete access to his personal health information, including contextual data such as his current location, to his family members.

Once a week, Ravi records his weight, blood glucose and blood pressure, using devices that send the measurements wirelessly to his mobile phone. At least once per day, Ravi enters his dietary information manually into his mobile phone. Due to his participation in the management program, Ravi's insurance company offers to reduce his premium if he shows significant improvement in controlling his diabetes. To demonstrate improvement, Ravi must provide the insurance company access to his health data. Ravi instructs his PHR to provide aggregate information of his activity, diet and physiological parameters to the insurance company. Ravi allows his doctor greater access to his PHR, however; using the PHR data, Ravi's doctor evaluates the impact of the diabetes-management program on Ravi's health, setting goals that he can achieve and improve his quality of life. (Scenario adapted from Kotz et al. [2009].)

We return to this scenario in Section 4 to explore the privacy issues in more detail.

2. SURVEY OF LEGAL AND PRIVACY FRAMEWORKS FOR HEALTHCARE IT

In this section, we first survey literature that describes existing legal frameworks, including laws and policy issues, related to healthcare information management and exchange. Next, we survey frameworks designed to identify privacy principles for healthcare information management systems.

2.1. Privacy Laws and Regulations (US)

Law and industry-wide policy should be the foundation for Patients' trust in the privacy and confidentiality provided by healthcare information systems. Unfortunately, most countries have only begun to think deeply about the policy issues and to pass thoughtful, meaningful laws. The US passed the Health Insurance Portability and Accountability Act (HIPAA) in 1996 [HIPAA 2010], intended to ensure portability and accountability across healthcare providers and insurers, and to identify important privacy rights and security policies. Unfortunately, gaps and weaknesses in the law left many critical consumer issues unresolved [Goldman 1998]. The American Recovery and Reinvestment Act (ARRA), passed in February 2009, extends HIPAA and fills many of those gaps [Steinbrook 2009; ACLU 2009; AMA 2009; HL 2009]. Furthermore, it provides the initiative for broad deployment of Electronic Health Records (EHR) by 2014.

The US Department of Health and Human Services (HHS) has funded prototypes of health information exchanges and chartered a federal advisory body (the American Health Information Community—AHIC) to make recommendations for accelerating the adoption of nationwide EHRs. Public/private collaborations in many states are developing linked health networks to share information among patients, providers and payers. (Providers are those entities that provide medical care; payers are those entities that pay for the care.)

Although several organizations have recommended policies and best practices for privacy and confidentiality in healthcare, and continue to do so, these efforts remain fragmented. Patients desire privacy, but the laws and policies that provide them control over their health information vary across providers and regions, and real mechanisms for control are lacking. Organizations and patients both consider confidentiality a key component in the doctor-patient relationship, yet there are disagreements about how one should implement broader privacy rights for patients when sharing information.

We expect that patients will be reluctant to accurately and honestly disclose personal information or seek medical care unless they can trust that their personal, sensitive information will be handled with the required level of confidentiality. Furthermore, we expect they will be reluctant to use mHealth devices if they feel that the laws and policies governing those organizations that deploy the devices, or collect the data, will not respect their privacy.

The best-known set of rules regarding privacy in healthcare arose from the HIPAA Privacy and Security Rules [HIPAA 2010]. HIPAA seeks to ensure (among other things) that personal health information in the possession of healthcare providers and payers is protected from uses and disclosures that would compromise interests of patients.

The HIPAA Privacy Rule provides the baseline of privacy protection for health information in the US, while allowing more protective state laws to continue in force. The Privacy Rule establishes a federal mandate for individual rights in health information, imposes restrictions on uses and disclosures of individually identifiable health information, and provides civil and criminal penalties for violations. The Privacy Rule of HIPAA addresses the use and disclosure of a patient's "protected health information" (PrHI) by healthcare plans, medical providers, and clearinghouses, also called "covered entities." The covered entities, which usually have direct interaction with the patient, are the primary agents capturing a patient's health information for a variety of purposes including treatment, payment, managing healthcare operations, or medical research. The Privacy rule of HIPAA includes standards for protection of health information in electronic form and requires covered entities to ensure implementation of administrative safeguards in the form of policies, personnel, physical and technical safeguards to protect, monitor and control intra- and inter-organizational access. As a concrete example from a major hospital, see how the Mary Hitchcock Memorial Hospital and Dartmouth-Hitchcock Clinics [2009] applies these regulations.

Individuals have certain key rights under HIPAA. These include the right to access, inspect and request a copy of their PrHI that covered entities hold, the right to request corrections to PrHI, the right to receive notice of how covered entities use their PrHI, the right to choose whether to permit use or disclosure of their PrHI, and the right to request an audit trail of PrHI disclosures.

HIPAA covers only those entities that provide healthcare (e.g., hospitals, clinics and doctors) or handle payment (e.g., insurers), but not other entities such as business associates that perform a variety of services (e.g., claims management or accounting) for hospitals and doctors, or PHR service providers such as Google Health or Microsoft HealthVault. Covered entities that hold PrHI may use it without an individual's consent for the purposes of providing treatment to the individual. However, the patient's signed authorization is required for using his PHI for other activities such as marketing, research, or fundraising. The Privacy Rule requires healthcare providers and insurers to obtain additional documentation from researchers before disclosing personal health information for research and to scrutinize researchers' requests for access to health information. Authorization should require either a written consent or Institutional Review Board (IRB) approval for waiver of authorization.

The HIPAA privacy rules are complex regulations expressed in legal English; it remains a challenge to interpret these rules and to implement them in clinical workflow and information systems. Several researchers have developed formal methods for extracting precise rules from the regulatory language, in ways that can guide software engineers implementing such systems [Barth et al. 2006; Breaux and Antón 2008, e.g.].

The American Recovery and Reinvestment Act of 2009 (ARRA) encourages the adoption of electronic medical records (EMR) by doctors and hospitals and imposes additional privacy safeguards [AMA 2009; HL 2009; ACLU 2009; CDT 2009]. Under ARRA, patients have the right to obtain an electronic copy of their EMRs and to

have it transmitted securely to a healthcare provider of their choice. ARRA expands the set of business associates that are "covered entities"; it remains unclear whether PHR providers are also covered entities. ARRA prohibits the unauthorized sale of medical records; it also requires covered entities to maintain an audit trail of accesses to patient data and notify patients of security breaches if their data is affected. ARRA allows patients paying out-of-pocket to restrict disclosure of their health information and imposes stricter penalties for violations of its regulations.

As mHealth comes into common use outside the clinical context, however, the legal protections of HIPAA and HITECH may not apply. It remains unclear, for example, whether mHealth records stored on a mobile phone, or in a PHR, may be subject to inspection by law-enforcement agents who obtain a search warrant. Suppose your PHR contains twenty years of location data, associated with your heart rate, skin conductance, and conversational state. This information would be interesting to a forensics lab investigating a crime: were you at the site of the crime? were there others present? was your heart racing? Similarly, this information would be of great interest to corporate marketing departments. What laws cover this kind of information, and protect individual privacy?

2.2. Privacy Laws and Regulations (Elsewhere)

There is no European law or directive specifically for privacy of health information. Most European countries have broad data-protection and privacy laws. These laws cover the gamut of personally identifiable data, which includes health-related information as a subset. The laws deal with privacy aspects such as legitimacy of need-to-know, notification, consent, individual's rights with respect to access, examination and amendment of data, and requirements for security safeguards; they also provide remedies and sanctions against violations. The laws are administered through independent national "data protection commissions" or "registrars" [Lowrance 2009]. In the UK, for example, the National Health Service (NHS) is building out a national-scale system of electronic health records [NHS 2009a], basing its privacy practices on the EU Data Protection Directive, described below.

The Organization for Economic Cooperation and Development (OECD) created a comprehensive data protection system across Europe [OECD 1980] based on the following privacy principles: data purpose specification, openness, individual participation, collection limitation, data quality, use limitation, data security safeguards, and accountability. Later, the EU Data Protection Directive (also known as Directive 95/46/EC) [EU 2009], which went into effect in 1998, became an important component of EU privacy law. The directive regulates processing of personal data and free movement of such data regardless of whether it is automated. Processing covers any set of operations on personal data such as collection, storage, modification, deletion, retrieval, or transmission. Broadly, it states that personal data can be processed only when certain conditions are met. These conditions fall in three categories: (a) when the individual has given informed consent or when processing is in the public interest (e.g., for law enforcement), (b) when the purpose is legitimate (as defined by the law), (c) when the data processed is relevant and not excessive to the purpose for which it was collected and is usually processed. It further states that personal data can be transferred to countries outside the EU only if that country offers adequate level of protection.

To bridge the differences in the US and EU privacy approaches and to provide a streamlined means for US organizations to comply with the EU Directive, the US-EU Safe Harbor Framework [Safe 2010] was developed. To opt-in to the program, US companies must adhere to the seven Safe Harbor principles related to notice, choice, onward transfer, security, data integrity, data access and enforcement. Organizations

can self-certify their compliance to these principles under the oversight of Federal Trade Commission.

In New Zealand, the 1993 Privacy Act [NZPA 1993] is one of the most comprehensive privacy acts outside Europe; it applies to handling of all personal information collected by businesses and governmental agencies. The legislation encompasses twelve principles based on OECD privacy guidelines and the National Information Principles contained in the Australian Privacy Act. Some important principles in the Act that are not part of OECD include: data collectors should specify the manner of collection, agencies should not keep information longer than necessary, and information should be checked for correctness before use. The 1994 Health Information Privacy Code [NZHIPC 2008] describes a code of practice for handling health information covered by the New Zealand privacy act.

In the Asia Pacific region, countries such as Japan, Hong Kong, Australia, South Korea, and Taiwan have adopted privacy legislation, whereas, Thailand, China, Malaysia, Philippines are still in the process of drafting privacy laws. Singapore [Kuie 2003], Vietnam and Indonesia have privacy laws covering only certain sectors. In Singapore, health information is protected by the common law, code of practice and Computer Misuse Act. India also does not have comprehensive privacy laws in place. Indeed, in India there is no law guaranteeing general right to privacy, though some elements of it occur in the common law, criminal law and Article 21 of the constitution upholding personal liberty [Brahmbhatt 2010]. The Information Technology Act of 2000 (and an amendment in 2008) loosely address certain aspects of data privacy, but without a precise definition of the "sensitive personal data or information" it intends to cover [Vadehra 2011; DIT 2011]. The Government of India released a new set of rules in April 2011, defining "sensitive personal data or information" (including medical records) and detailing rules about their collection, storage, and transfer; the rules are largely consistent with OECD guidelines [India 2011].

2.3. Law and mHealth

To our knowledge, none of the laws or regulations described above specifically address *mobile* health issues. In the US, most of the laws and regulations that apply to healthcare are focused on the clinical setting, and since many mHealth applications we imagine occur outside the clinical setting, it may be that such mHealth data are not protected by current regulations. In the EU, most of the laws and regulations apply to personally identifiable data, regardless of the application or setting, and thus they may apply to mHealth data and to all the various uses of mHealth technology. A legal analysis is outside the scope of this article.

Furthermore, because mHealth devices are mobile, a traveler may encounter a variety of legal frameworks. It may be necessary to develop international legal standards to protect travelers who choose to (or need to) wear mHealth devices in countries with lower standards of privacy than their home country.

One of our concerns, as mHealth proliferates, is that the data is often collected and stored by the device vendor or, in some cases, by the network operator (e.g., in closed-model cellphones where the operator controls the placement of applications on the phone). It seems likely that, at least in the US, mobile operators may have access to health data but not be covered by current laws that regulate PHI. Although these problems may also arise with other kinds of data, such as financial data, it is greatly concerning that the Patient may have little knowledge or control over the flow of mHealth sensor data and its secondary uses by device vendors, platform providers, and network operators.

2.4. Conceptual Privacy Frameworks for Healthcare

Although national laws lay the foundation for health-data privacy, in some countries, as technologists we seek to ground our systems in fundamental principles that protect the privacy of the Patients using mHealth systems. We begin by developing a conceptual privacy framework that lays out a coherent set of principles; fortunately, we can draw on recent frameworks developed by healthcare groups. We define a *conceptual privacy framework* to be a coherent set of actionable principles to protect Patients' health information privacy. When developing and deploying a healthcare information system, the system design should include security and privacy properties that align with the principles in the conceptual privacy framework.

In the next section, we propose such a framework, and a corresponding set of privacy properties, for use in the following sections to survey the technology literature and to identify important research questions. As groundwork, we begin by distilling the essential privacy principles from frameworks produced by four key entities in the US—the Office of the National Coordinator, the Health Privacy Project, the Markle Foundation, and the Certification Commission for Healthcare Information Technology.

ONC. In 2008, the Office of the National Coordinator for Health Information Technology (in the US Department of Health and Human Services) announced its *Nationwide Privacy and Security Framework for Electronic Exchange of Individually Identifiable Health Information* [ONC 2008]. This document, which we call the "ONC Framework", provides the newest and most authoritative privacy framework for healthcare released in the US.

BP. In 2007 the Health Privacy Project listed 10 "best practices" for employers who are developing personal health records (PHR) [HPP 2007]; we refer to these best practices with the abbreviation BP. (The Health Privacy Project has recently been adopted by the Center for Democracy & Technology, CDT.)

CF. The Markle Foundation's project "Connecting for Health", which brings together a wide range of stakeholders, developed a "Common Framework" as a model for healthcare information exchange [MF 2008]. The Connecting for Health Common Framework (CF) describes both policy and technical principles for healthcare information exchange, and provides concrete prototypes of each: everything from patient consent forms to data-interchange formats and information architecture. The Center for Democracy & Technology later endorsed the policy aspects of the Common Framework in their own policy document [CDT 2008].

CCHIT. The Certification Commission for Healthcare Information Technology (CCHIT) is a nonprofit organization that certifies healthcare information systems, and they recently released their certification criteria for PHRs. Although it appears that they have a process in place to determine these criteria, and that the process may enable and encourage updates to the criteria, we refer the list of criteria as published in a booklet copyrighted in 2008 [CCHIT 2008].

In this article, we refer to each principle using labels that reflect the order in which it appears in the source document. For example, the INDIVIDUAL ACCESS principle appears as the first item in the ONC framework and the fifth item in the Markle Foundation's Common Framework (CF); therefore, we refer to this principle via labels ONC1 and CF5, respectively. Appendix A (or, equivalently, a workshop paper [Kotz et al. 2009]) provides a labeled list of the principles in each of those four frameworks, and some commentary on other frameworks.

For brevity, we summarize the principles in the existing frameworks, making reference to the original principles.

Openness and Transparency. Healthcare information policies, procedures and technologies should be open and transparent to Patients who use the healthcare information system. ONC3 narrows the scope of the policies, procedures and technologies to those that directly affect Patients and/or their individually identifiable health information. BP1 requires employers that offer PHRs to their employees to be transparent about their reasons for such offers, and about all policies that apply to the PHR and any updates to those policies. CF1 requires complete openness about information collected: what was collected, why (purpose of its use), who (can access and use it), where (it resides) and how (each individual may obtain access to it and control who has access to it).

Individual Participation, Access, Consent, and Control. This principle captures the idea that every individual must be able to control the privacy of his or her healthcare information. ONC1 focuses on individual access and suggests that individuals be provided with simple, timely and readable access to their individually identifiable health information. ONC4 focuses on individual choice and suggests that individuals be provided reasonable capability and opportunity to make informed decisions about the collection, use, and disclosure of their individually identifiable health information. BP4 requires employees to control who is allowed to access their PHRs and prohibits employers from accessing a PHR to obtain employees' individually identifiable information; furthermore, employees should be able to choose whether or not to grant access to personal health information in their PHRs for any "secondary uses"; and, the employee should be able to view an audit trail that shows who has accessed his PHR data. CF4 suggests that individuals should be provided the ability to control access to their personal information, be provided with the knowledge of who is storing what information on them and how that information is being used and be provided the ability to review the way their information is being used or stored. CCHIT1 focuses on individual consent, requiring certified PHRs to include safeguards that either require an individual to provide explicit consent before her account is opened, or allow the individual to opt out of the service; it also requires the safeguards to allow an individual to decide if her data can be collected, displayed, accessed, stored, released or disclosed. CCHIT2 requires the PHR to enable an individual the ability to decide what information is private and to restrict access to it; furthermore, the PHR provider must obtain an individual's permission to gather or disseminate any information about her. CCHIT2 also suggests that individuals must decide who else can view information in their PHRs, and limit the types of information that can be viewed.

Purpose Specification and Data Limitations. This principle captures a foundational requirement in a healthcare information system that meets individuals' privacy requirements. According to this principle, each provider that needs to collect, store and use an individual's healthcare information must specify the purpose for which the information will be collected, stored and used; subsequently, that entity must limit collection, storage and use to only the specified purpose. If the purpose changes in the future, the entity must inform the individual of such a change and subsequent use. ONC5, CF2, CF3, CF4, and CCHIT3 describe this principle.

Data Quality and Integrity. This principle identifies the need to ensure that individuals' information—both individually identifiable information and health information—within the health information system is complete, accurate, relevant to purpose and up-to-date. The information must be recorded accurately at the time of collection and must continue to remain accurate even after any modifications to the system. Any changes to the recorded information must be appropriately annotated to reflect those changes accurately. The health information system must ensure that individuals' information cannot be altered or destroyed in an unauthorized manner. ONC6, BP8, and CF6 describe this principle in the respective documents. CCHIT does not explicitly list this principle.

Security Safeguards and Controls. This principle focuses on the security of the infrastructure upon which the health information system is built. It requires that the infrastructure provide appropriate mechanisms, including physical, technical and administrative mechanisms, to keep individuals' health and identity-related information confidential and prevent unauthorized access, use or disclosure of that information. This principle further requires the health information system to ensure availability of data to its users. ONC7 and CF7 describe this principle in general. BP7 explicitly requires employer-provided PHR systems to have robust authentication mechanisms for access to PHRs and also requires PHR systems to maintain an audit trail that captures activity associated with each PHR. CCHIT7 only focuses on data availability, requiring the PHR system to ensure that individuals' PHRs are available to them when needed.

Accountability, Oversight and Remedies. This principle captures the requirement that all entities in the health information system, including providers, payers, and employers, must be held accountable to ensure that they adhere to the above principles in letter and spirit. Monitoring mechanisms must be in place to capture adherence (or the lack thereof) to these principles. When inappropriate use, access or disclosure of information occurs, the entity or entities responsible for those actions must inform and appropriately compensate affected individuals and take appropriate remedial measures to prevent future occurrence of such incidents. ONC8 and CF8 describe this principle. BP9 explicitly lists the steps that employers should take if inappropriate access or use of information contained in an employee's PHR occurs. CCHIT does not describe this principle.

We build on these essential principles in the following section, in which we define a conceptual privacy framework for mHealth, and list specific privacy properties for mHealth systems.

3. AN MHEALTH PRIVACY FRAMEWORK

So, after reviewing all those conceptual privacy frameworks, which one do we recommend as the basis for research and development in mHealth systems? Both the ONC Framework and the Common Framework (CF) are appealing, because both are recent (2008), both are fairly complete, and both were developed by diverse groups of experts and stakeholders. We chose to ground our recommendations on the Common Framework for four main reasons.

First, the CF more clearly states the fundamental privacy principles for healthcare information systems. The ONC Framework leaves many important issues to the details, rather than expressing them in the main bullets. Indeed, the ONC Framework leaves certain issues implicit. When privacy is at stake, it is important to be explicit. Second, the ONC Framework is less concrete, leaving flexibility on several principles. We believe it is important to state the core principles, clearly and simply, and to aim implementations at them. Third, the CF has a more Patient-centric viewpoint. Finally, concrete materials accompany the CF principles: everything from patient consent forms to data-interchange formats and information architecture.

We thus build our mHealth Privacy Framework by adopting all of the Common Framework and adding complementary pieces of the other frameworks; we add one mHealth-specific principle. In the list that follows, we quote from or refer to other frameworks; the bracketed codes (e.g., ONC3) refer to the labels in Appendix A.

MH1. *Openness and Transparency.* "Consumers should be able to know what information has been collected about them, the purpose of its use, who can access and use it, and where it resides. They should also be informed about how they may obtain access to information collected about them and how they may control who has access to it" [CF1]. The system should be open about the policies and technologies in use [ONC3].

MH2. *Purpose Specification*. "The purposes for which personal data are collected should be specified at the time of collection, and the subsequent use should be limited to those purposes, or others that are specified on each occasion of change of purpose" [CF2].

MH3. *Collection Limitation and Data Minimization*. "Personal health information should only be collected for specified purposes and should be obtained by lawful and fair means. The collection and storage of personal health data should be limited to that information necessary to carry out the specified purpose. Where possible, consumers should have the knowledge of or provide consent for collection of their personal health information" [CF3].

MH4. *Use Limitation (Transitive)*. "Personal data should not be disclosed, made available, or otherwise used for purposes other than those specified" [CF4]. "The information policies and practices . . . should follow the data through chain of trust agreements that require business partners to adhere to the . . . applicable policies and practices" [BP6].

MH5. *Individual Participation and Control*. "Consumers should be able to control access to their personal information. They should know who is storing what information on them, and how that information is being used. They should also be able to review the way their information is being used or stored" [CF5]. Patients should be able to make informed choices about what data is collected, how it is used, and to whom it is disclosed [ONC4]. Patients "can designate proxies to act on their behalf" [BP5].

MH6. *Data Quality and Integrity*. "All personal data collected should be relevant to the purposes for which they are to be used and should be accurate, complete, and up-to-date" [CF6]. Patients should be able to correct mistakes in their records [ONC2].

MH7. *Security Safeguards and Controls*. "Reasonable safeguards should protect personal data against such risks as loss or unauthorized access, use, destruction, modification, or disclosure" [CF7].

MH8. *Accountability and Remedies*. "Entities in control of personal health information must be held accountable for implementing these principles" [CF8]. "Remedies must exist to address security breaches or privacy violations" [CF9].

MH9. *Patient Access to Data*. Patients should have an easy method to obtain their PHI in a readable electronic format [ONC1]. Patients "should be able to annotate the records submitted by others, as well as to enter their own information, with [Patient]-entered data marked as such" [BP3].

MH10. *Anonymity of Presence*. The presence of medical sensing devices, or the nature of sensor-data collection, should not be observable by nearby parties (this privacy threat is unique to mHealth).

These principles provide broad guidelines for an mHealth system. Most of them derive from principles developed for a clinical healthcare setting, where the context is a Patient's medical record; when applying them to mHealth one must remember that mHealth systems may collect more data, including personal activities outside the clinical setting, over months or years, and thus require a more thoughtful interpretation of issues like purpose specification (MH2), collection limitation (MH3), or use limitation (MH4).

We next derive a specific list of properties needed in an mHealth system that follows the above principles; these properties are more directly useful to, say, anyone designing or evaluating an mHealth system.

3.1. mHealth System Properties

This mHealth privacy framework provides us a foundation to identify a set of privacy properties that a high-quality mHealth system must implement. Again, the codes refer

to these labels (e.g., MH1) or Appendix A (e.g., ONC3). Finally, for completeness we list the functional properties, including integrity, availability, and auditability properties.

3.1.1. Security and Privacy Properties. A privacy-aware mHealth system should provide the following.

P1. Inform Patients (MH1, MH2, CF1, CF2, ONC1, ONC4)
  a. *What* PHI is collected and stored
  b. *Why* PHI is collected and stored
  c. *Where* PHI is stored and at which organization
  d. *Who* has access to their PHI, and under what circumstances
  e. *When* PHI collection purpose (why) changes or access (who) changes
  f. *How* their PHI is used
  g. About *risks* of data collection or disclosure
  h. About security *breaches* or PHI misuse
P2. Enable Patients to review storage and use of their PHI (MH1, MH5, CF5)
  a. Review historical records of all information in property P1.
P3. Enable Patients to control, through informed consent (MH1, MH3, MH5, CF1, CF3, CF5, ONC4, BP5),
  a. What PHI will be collected and stored, and in what contexts,
  b. When PHI will be collected and stored (allowing Patients to stop and restart data collection),
  c. Who will have access to their PHI (including Patient proxies), and in what context, and
  d. How their PHI may be used, and in what circumstances
P4. Provide access to PHI
  a. Enable Patients to access their PHI (MH9, CF1, ONC1, BP3).
  b. Honor Patients' requests to add, annotate, correct and delete their PHI (MH6, MH9, CF6, ONC2, BP3), wherever possible.
P5. Provide easy-to-use interfaces for all of the above. For example, an interface for ensuring informed consent might highlight keywords in the provider's privacy policy and display their definitions using mechanisms such as a tooltip.
P6. Limit collection and storage of PHI (MH2, MH3, CF2, CF3)
  a. As needed for specified purpose
  b. Per limitations of Patient consent
  c. Using lawful and fair means
P7. Limit use and disclosure of PHI to those purposes previously specified and consented (MH2, MH3, MH4, CF2, CF3, CF4)
  a. Policies should follow PHI as it flows to other entities (MH4, BP6)
P8. Ensure quality of PHI (MH6, CF6)
  a. Ensure data freshness and accuracy during collection
  b. Ensure data integrity and completeness during transmission, processing, and storage
  c. Ensure authenticity of Patient providing input or wearing sensor
  d. Ensure authenticity and quality of sensors
P9. Hide Patient identity, sensor presence and data-collection activity from unauthorized observers (MH10)
P10. Support accountability through robust mechanisms (MH8, CF8)
  a. Include audit logs for all access, addition, deletion, and modification of PHI (the MN, too, should log its actions with respect to collection, upload, and access to PHI, and pairing with SNs and other devices)
P11. Support mechanisms to remedy effects of security breaches or privacy violations (MH8, CF9)

3.1.2. Functional Properties. A high-quality mHealth system should also have these properties.

P12. *Flexible*, supporting multiple types of data
  a. Streaming data, that is, high-frequency, periodic data
  b. Event data, that is, low-frequency aperiodic data
  c. Patient-entered data, using one or more modes of input
P13. *Scalable*, to large numbers of participants (Patients and Consumers) and devices (MNs, SNs)
P14. *Efficient*, particularly resources on SN & MN (memory, bandwidth, energy)
P15. *Usable* by
  a. Patient: physical usability of sensors, that is, preserve wearableness
  b. Patient and Provider: easy interfaces for data collection and access
  c. Physically challenged Patient: accessible interfaces for informed consent, and control over PHI
P16. *Manageable*
  a. Ensure remote configurability and calibration of system components
  b. Ensure ability to remotely manage lifecycle of software and credentials in system despite having no control over OS and firmware updates of system machinery
  c. Enable easy provisioning and de-provisioning of health applications, tools and data
P17. *Available*, preventing loss of (or loss of access to) PHI data
  a. At MN, data latency and risk of loss must balance against resource limitations
P18. *Interoperable*, across devices and time:
  a. multiple classes and brands of SN must interoperate with
  b. multiple classes and brands of MN, and must upload to
  c. multiple HRS from multiple vendors;
  d. must support legacy devices, and work with data from decades past.

We recognize that there are subtle and, in many instances, challenging aspects to supporting these properties in an mHealth system, and that many of these properties are not unique to mHealth. Some of these properties will require support on the SN and MN; others may be achieved only in the back-end servers or the data consumer's system. Careful study is needed, in the context of a system design, to identify which properties will need support on the mobile platform. For example, the interface designed to ensure property P1 (i.e., Inform the Patient) on a mobile node must be highly user-friendly because of the limited viewing area on its screen; a multimodal interface may best convey information to the Patient. As another example, because mobile nodes are highly vulnerable to physical loss or theft, they must be remotely manageable; that is, an administrative entity must be able to disable or lock the mobile node remotely and prevent Patient data from being stolen.

We next present a case study that provides specific examples to motivate the above properties.

4. CASE STUDY

Here we recall the "Diabetes Management" scenario (Section 1.4) to illustrate how some of the aforesaid properties can be realized in a "privacy aware" mHealth system.

Ravi is a diabetic who finds it difficult to manage his condition effectively, resulting in significant variation of his diurnal blood-glucose levels, and frequently elevated blood pressure and cholesterol levels. Ravi's doctor advises him to subscribe to a Diabetes Management Program offered by his hospital. As part of the program, Ravi wears a hospital-provided device that continuously monitors his activity level and calories

Fig. 1.   Screenshot of example policy.

burned. The device is designed as a wrist watch to improve usability and to prevent anyone in Ravi's vicinity from detecting that he is wearing it (P9–hide sensor presence). Upon first use, the device records features of Ravi's pulse waveform (plethysmograph). Whenever Ravi takes the device off his wrist and wears it again, the device uses the plethysmograph as a biometric to authenticate him. As long as Ravi wears the device, it monitors his activity and wirelessly sends encrypted data to his smart phone.

Also as part of the program, Ravi's smart phone is provisioned with certified software and a set of cryptographic keys. Upon installation and setup, the software on the smart phone uses one or more of the keys to authenticate communications with the wrist device. The software further verifies certificates claiming that the device is properly calibrated and untampered, thereby ensuring data quality and integrity (P8–ensure quality of PHI). The software processes data it receives from the device and acquires other contextual information such as Ravi's location (using GPS or Wi-Fi localization), ambient audio (using a microphone), calendar schedule and time of the day. This data is used to infer social context (such as whether Ravi is at home or at the office, in conversation or alone). The software reminds Ravi to take his medication, alerts him to long periods of inactivity, encourages him to log diet information and tracks his daily goals of calorie intake and expenditure. The phone synchronizes data with Ravi's PHR server.

Before starting data collection, the smart phone application presents Ravi with a privacy policy; one screen of a sample policy is shown in Figure 1. Ravi observes that the privacy policy is easy to understand, clear and concise (P5–easy to use interfaces), unlike verbose privacy policies on some health websites. The privacy policy clearly conveys (P1–inform patients) information about the collection, use, and storage of Ravi's Personally Identifiable Information (PII, a super-set of PHI).

After disclosing information about what PII will be collected and why, who has access to the PII and how it will be used, the software seeks Ravi's consent to begin using the system (P3–informed consent). Specifically, the software presents Ravi options in the form of check boxes that allow him to control data collection and dissemination. It allows Ravi to control data collection by enabling him to specify times of the day during which data can be collected by turning off location sensing, audio sensing, and medical

sensing (P3–enable patient to control data collection and dissemination). It also allows Ravi to control data dissemination by specifying who (people, roles, or organizations) has access to his PHI and at what level of detail. Ravi allows detailed access to his PHI (including context) to his doctor and spouse, allows the insurance company access to his aggregate health trends, and allows researchers access to anonymized aggregate data only. Entities that access Ravi's data are presented a policy (derived from Ravi's preferences) describing the reasonable and permitted usages of Ravi's PHI (P7–Policy should follow data). Ravi also names his spouse as a "proxy," allowing her to act on his behalf and with all privileges equivalent to his. These preferences are configured by Ravi in the hospital under the supervision of an expert who can answer Ravi's questions about the privacy issues and policy interface, but Ravi can revisit this interface and change his preferences at any time. The PHI software also allows Ravi to monitor where his PHI data resides (with which entities, at what locations) and for what purpose (P3– review storage and use of PHI).

An embedded accelerometer in the wrist device monitors Ravi's activity in terms of number of walking steps, and an optical sensor monitors his pulse rate (which is used to estimate exercise intensity). The system is able to accommodate real-time streaming data, periodic and aperiodic sampling of location information, calendar events, and audio; it also allows Ravi to manually enter diet information (P12–flexible). A fusion algorithm on the smart phone estimates Ravi's context, which serves both to augment the medical data in the PHR and more immediately to provide customized and relevant reminders, tips, and motivational messages that encourage him to increase physical activity and control his diet. For example, when the smart phone determines that Ravi is running late for a meeting, it avoids prompting him to take the stairs instead of the elevator. Or, when the system determines that Ravi is in conversation, it avoids using spoken motivational messages. Despite having access to multiple health and context sensors, the system limits raw data collection only to health sensors. In contrast, only the high level context inferred from the context sensors (such as microphone) is stored (P6–Limit data collection and storage).

Using authentication keys retrieved from a secure internal storage area, the smart phone periodically connects with the backend server to upload his activity, diet and context information, and to obtain messages for Ravi that arise from either automated or human analysis of his data. Ravi also receives a username and password for a secure web interface accessible from any computer (P4–enable patients to access, add and annotate their PHI). The web interface presents detailed charts and trends of his health parameters, activity, diet information and established context, which help him to review, introspect and improve his lifestyle. Per Ravi's privacy choices, the same information is available to his spouse and doctor for review. However, the insurance company may only access high-level information about his fitness level. This arrangement is sufficient for the insurance company as Ravi's insurance premium depends on his fitness level.

Whenever Ravi is wearing the wrist device, he can seamlessly use his smart phone to view the backend-stored PII. When the wrist device is not present, or cannot biometrically verify Ravi as its wearer, the smart phone will only provide access to PII after entry of the password. This approach provides ease of use in the common case but prevents misuse in the case of device loss or device theft.

The backend server maintains audit logs (P10–support accountability) of all accesses, additions, and updates to Ravi's PII. With his password, Ravi can review at any time who has accessed which parts of his PII and when. The system also has mechanisms to send an alert SMS to Ravi's smart phone in case there is an unauthorized access to Ravi's PII or if there is a security breach. In case of unauthorized access or security breach, Ravi can disable all the previously granted access permissions and freeze

his PHI account by calling the hospital's call center (P11–remedy effects of security breach).

Periodically, Ravi's smart phone receives mHealth software updates, because the software comes with a maintenance contract. Security-related patches are automatically installed, and new features are optionally installed. This remote manageability (P16–manageable) protects Ravi's PII and maintains Ravi's interest with the addition of novel features.

After one month, Ravi's smart phone is stolen while he is at the gym. With one call to his cellular carrier, his old phone is remotely disabled (destroying all the keys and PII stored on the phone) and a new phone is sent by courier. Fortunately, Ravi's service contract (with the mHealth software provider) also supports him in the case of a lost device, so the new phone arrives preconfigured with the necessary software and encryption keys, and automatically recognizes the sensors Ravi had paired with his earlier phone.

5. INFORMATION TECHNOLOGY IN HEALTHCARE

In this section, we begin with a look at privacy issues in healthcare information technology and in mobile-sensing technology, with an emphasis on the threats to privacy in the mHealth context. We end with a discussion of human-interface challenges and the sociological aspects of privacy management.

5.1. Privacy Technology for Healthcare

Recalling the NCVHS definition of privacy (in a healthcare setting) as the user's right to "control the acquisition, uses, or disclosures of his or her identifiable health data" [Cohn 2006], a threat to user privacy is the possibility that his right to control his PHI is weakened or eliminated due to erroneous or malicious actions. When these threats are realized, the consequences can be severe: exposure of identifiable Patient health data leading to loss of money or reputation, time spent recovering from medical identity theft, harm to health, or even death.

Table I summarizes these threats, organized by the type of threat: misuse of Patient identities, unauthorized access or modification of PHI, or disclosure of PHI. For each category, we consider three types of adversary: the *Patient* himself or herself, *insiders* (authorized PHR users, staff of the PHR organization, or staff of other mHealth support systems), and *outsiders* (third parties who act without authorization).

In the following sections, we survey existing technological approaches to mitigate these and related threats.[2] We draw on the literature in healthcare information technology, mobile computing, pervasive computing, wireless networks, sensor networks, cryptography, and computer security.

5.1.1. Identity Threats. In the first section of Table I we explore threats related to Patient identity. There are three concerns here. First, the Patient may lose (or share) their identity credentials, enabling others to have access to their PHI in the PHR (or perhaps PHI in their MN). The result may be a reduction of Patient privacy, because others may read, modify, or disclose their PHI (affecting privacy properties P3, P6–P9). In addition, insiders may use Patient identities for medical fraud, for example, by submitting fraudulent insurance claims [Dixon 2006]; the result can be financially or even medically damaging to the Patient (affecting privacy properties P7, P8). Furthermore, in the growing problem of medical identity theft, outsiders (or insiders) may use a Patient's identity to obtain medical services [Johnson 2009], potentially with financial or medical damage to the Patient (affecting P7, P8). Finally, in some settings (such as

---

[2]A shortened version of this section appeared as a workshop paper [Kotz 2011].

Table I. Privacy-Related Threats in mHealth Systems.

Identity threats: misuse of patient identities

patients	leave PHR credentials on public computer (identity loss)
patients	share passwords with outsiders (identity sharing)
patients	reveal passwords to outsiders (social-engineering attack)
insiders	misuse identities to obtain reimbursement (insurance fraud) [Dixon 2006]
insiders	misuse identities to obtain medical services (identity theft) [Johnson 2009]
outsiders	misuse identities to obtain medical services (identity theft) [Johnson 2009]
outsiders	reidentifying PHI in de-identified data sets [Sweeney 2002; Malin 2006]
outsiders	observe patient identity or location from communications

Access threats: unauthorized access to PHI or PHR

patients	consent preferences, as expressed, do not match those desired
patients	intentional (or unintentional) access beyond authorized limit
patients	mistaken modifications, because of over-privilege or inadequate controls
insiders	mistaken modifications, because of over-privilege or inadequate controls [Sinclair and Smith 2008]
insiders	intentional unauthorized access, for curiosity or malice [Appari and Johnson 2010; Sinclair and Smith 2008]
insiders	intentional modifications, to obtain reimbursement (insurance fraud) [Dixon 2006]
outsiders	intentional unauthorized access, for curiosity or malice [Messmer 2008]
outsiders	intentional modifications, for fraud or malice [Messmer 2008]

Disclosure threats: unauthorized disclosure of PII and PHI

data at rest, in the PHR:	
patients	inadvertent disclosure due to malware or file-sharing tools [Johnson 2009]
insiders	inadvertent disclosure due to malware or file-sharing tools [Johnson 2009]
insiders	inadvertent disclosure due to sharing passwords [Sinclair and Smith 2008]
insiders	intentional disclosure, for profit or malice [Appari and Johnson 2010]
outsiders	intentional disclosure, for profit or malice [Appari and Johnson 2010]
data at rest, in the mobile devices:	
patients	loss of MN or SN exposes PHI, keys, SN types, sensing tasks
outsiders	theft of MN or SN exposes PHI, keys, SN types, sensing tasks
data in transit:	
outsiders	eavesdrop on SN-MN, MN-PHR, PHR-PHR, PHR-client; traffic analysis and/or content decryption [Srinivasan et al. 2008, for example]
outsiders	observe presence and type of sensors on patient [Halperin et al. 2008b]

research) Patient identities are removed from the PHI, and the risk is that an outsider may combine the de-identified data with data from another source to reidentify the Patients, that is, to relink Patient identity to their PHI [Sweeney 2002; Malin 2006] (affecting P7).

In this section we survey work related to authentication, anonymization and re-identification, and location privacy.

Authentication. Authentication protocols and mechanisms are used to authenticate the Patient (to ensure that the correct Patient is being sensed), to authenticate the provider (to ensure that only authorized personnel have access to the medical equipment or sensor data), and to authenticate devices (to ensure that only valid sensing equipment can participate, and that data is sent to the authentic information systems). Authentication is a foundational technology that impacts all our privacy properties, because without authenticating the Patient it is impossible to provide the Patient correct information and controls (P1–P4). Authentication failures expose PHI to disclosure or modification (P6–P9) and subvert proper accounting (P10), which makes it more difficult to remedy breaches (P11). Poor authentication interfaces (P5) can be particularly troublesome, because they encourage unsafe user behavior (such as password sharing).

Authenticating the Patient. The most common method of authenticating Patients to a PHR or other healthcare IT system is to verify a combination of their username and

password. Of course, this method is susceptible to a variety of well-known attacks. Some PHR providers are testing or deploying two-factor authentication [Moore 2009]; for example, PHRAnywhere, a payer-based PHR provider, uses smart cards to authenticate Patients. Users of Microsoft's HealthVault can login to their PHR accounts using their Hotmail or Windows Live identities (which are based on username/password) or with OpenID accounts (many of which offer a second factor of authentication via physical USB keys).

In mHealth applications there are additional challenges. The data consumers need to be sure that the data collected by sensors is collected from the correct Patient; otherwise the data may be interpreted incorrectly and result in treatment errors or incorrect conclusions in research or public-health monitoring [Cornelius and Kotz 2010]. The sensors may be attached to the wrong person by mistake—for example, in a hospital setting where the nurse confuses two identical sensors, each previously configured for a particular Patient, and places them on the wrong Patient. In a home-health setting, the Patient may not realize that it is important to wear his sensor rather than the identical one assigned to his wife. Finally, in some scenarios, the Patient may be motivated to cheat by attaching the sensor to a healthier friend or relative. Cornelius and Kotz [2011] developed a solution to this problem by including an accelerometer in each wearable sensor, and correlating accelerometer readings across devices.

There is extensive research related to the use of biometric data for authentication, or more correctly, for identity verification [Cherukuri et al. 2003; Bekiaris et al. 2008; Jain et al. 2007, 2004]. Some biometrics are not sufficiently unique to be used for identification, that is, to identify the subject given only the biometric signature. Instead, the subject claims a particular identity and the biometric data is used to verify whether they are likely to be the person associated with that identity. There are two important cases. First, one must verify the Patient's identity when the sensor node is first attached to the Patient, or associated with the Patient's mobile node. Second, in some settings it is necessary to periodically re-verify that the sensor remains attached to the Patient, and indeed to the same Patient. Several research studies propose methods based on features from electrocardiography (or similar biometrics) to verify Patient identity [Sriram et al. 2009a; Irvine et al. 2008; Jea et al. 2008; Agrafioti and Hatzinakos 2008]. In constrained settings, such as a household with small population, it is possible to distinguish residents by height [Srinivasan et al. 2010] or weight (in a commercial bathroom scale from Withings). It remains an open problem to find a robust biometric that is usable, inexpensive, and reliable under a range of Patient activities.

Authenticating the Provider. The HIPAA Security Rule (Section 164.312(d) Person or Entity Authentication) states that covered entities must "implement procedures to verify that a person or entity seeking access to electronic protected health information is the one claimed" [Scholl et al. 2008], and the 2009 HITECH Act extends this rule to business associates [HITECH1 2009; HITECH2 2009]. Most of the issues here are the same as for authenticating Patient access, above. The US Department of Veterans affairs has implemented "single sign-on" authentication to enable users to easily access multiple portals of the department using only one set of government-issued credentials.

The National Health Service (NHS), in the UK, is one of the largest national-scale EHR projects underway anywhere. The NHS Care Records Service (CRS) is a secure service that links patient information from different parts of the NHS, enabling authorized NHS staff and patients to have access to the health information [NHS 2009b]. The NHS CRS has two elements, a Detailed Record and a Summary Care Record. Patients' Detailed Records are securely shared between the local NHS entities, such as the doctor and a hospital, but only the summary of their health information is available to NHS staff at a national level.

The NHS CRS includes several data security and access mechanisms. A registration authority within each healthcare organization is responsible for verifying the identities of healthcare professionals and support staff, and to register all the caregivers allowed to access NHS CRS. The registration authority issues a smart card to each caregiver, which enables the caregiver to authenticate to CRS; the smart card is printed with the user's name, photo and a unique identity number. The system uses these cards and identities to provide role-based access to patient information. For example, a doctor's receptionist may access the information required for appointment scheduling, but not the details in the patients' clinical record. CRS maintains an audit trail of all accesses, which is available to the patient on request. However, there have been several reports of "inappropriate access" by sharing of passwords and PINs by the staff, which poses a serious insider risk [Collins 2006]. Reports have alleged that some doctors allowed their support staff to use their login IDs to save time (due to slow log-in performance), particularly for critical patients requiring emergency care [Collins 2006]. This experience highlights the importance of a balance between the robustness of security solutions and their usability.

Authenticating Devices. In our reference architecture, there are several component devices and services. In any secure solution, these components must be able to authenticate each other. When a mobile node (MN) communicates with sensor nodes (SNs), it must determine whether they are *authentic* sensors, that is, they are *valid* (truly the sensors they claim to be), *untampered* (not compromised by an adversary), and *correct* (they are the specific instances attached to the desired Patient, not another Patient nearby). Similarly, the SNs must determine whether the MN requesting data is the correct MN, that is, the one authorized to receive sensor data. Finally, MNs provide reports to, or obtain configuration from, healthcare services; these transactions also require mutual authentication so that the service believes the data comes from the correct MN (really, the correct Patient), and the MN believes the data is sent to (or configuration comes from) an authorized healthcare service.

Fundamentally, these authentication challenges may be easily solved by asymmetric cryptography and a public-key infrastructure. The problem is not so simple, however, for five reasons. First, these mobile devices are necessarily small and their resources are limited; asymmetric cryptography is computationally expensive. Second, these devices are often disconnected from the Internet, or have a weak connection to the Internet, obviating solutions that assume a live connection to (for example) a certificate authority. Third, these devices are small and may be easily lost or stolen, leading to a loss or exposure of embedded keys. Fourth, key distribution and key storage require secure and *easy-to-use* solutions, and yet some nodes have little or no human interface. Finally, some of the devices (notably the MN) may be owned and configured by the Patient rather than by the medical device manufacturer or healthcare provider.

We have seen few solutions in the literature that attempt to address the broad key-management challenge in the context of mHealth. One approach considers mote-class sensor nodes and demonstrates techniques for secure key exchange, biometric methods to authenticate the Patient, and an encryption protocol to protect the sensor data [Malasri and Wang 2008, 2007]. These two papers also cite several other papers related to the use of motes in mHealth applications.

Key distribution is a challenge in sensor networks [Xiao et al. 2007], but most existing solutions do not apply to the mHealth context, because they imagine a multi-hop network architecture, they imagine that the sensor nodes are homogeneous and indistinguishable, and the keys may be distributed for a particular purpose, such as authenticating software updates from the central server, encrypting sensor data for the central server, signing sensor data, or authenticating routing information; the proposed mechanisms are often subtly related to these purposes. Finally, they usually assume

that the set of sensor nodes is fixed; in many mHealth settings, sensors may come and go, or even be shared (as in a hospital where a nurse moves a blood-pressure cuff from Patient to Patient, or in a home where an elderly couple shares a bathroom scale).

Most sensor-network papers note, as conventional wisdom, that conventional cryptography is too challenging for limited-resource sensor devices. We note, however, that even mote-class devices can support encryption methods like ECC [Liu and Ning 2008] and Rijndael [Vitaletti and Palombizio 2007]. (One paper considers ECC on motes specifically in the healthcare setting [Mišić 2008].) Furthermore, recent work has demonstrated the viability of public-key encryption as well [Watro et al. 2004; Malan et al. 2008]. Others have demonstrated that it is feasible and inexpensive to couple mote-class sensor nodes with hardware-encryption support like that in a TPM [Hu et al. 2009]. Finally, hardware-accelerated crypto is increasingly common; it appears in cheap, microSD smart cards [GD 2011, for example].

Finally, since many mHealth scenarios imagine an often-disconnected environment with weak connectivity to the Internet, offline approaches to PKI are important. MNs and SNs should be able to store the public keys for a small but sufficient number of certificate authorities, and even to store large revocation lists—megabytes are cheap—but mHealth nodes may have neither the connectivity nor the bandwidth to keep those lists up to date. The topic of certificate revocation lists has a rich and controversial history [Rivest 1998; McDaniel and Rubin 2000; Gutmann 2002, for example], but we need to look at alternatives. Our mHealth scenarios need certificate-revocation solutions that require neither a revocation list nor strong Internet connectivity for online queries to the certificate authority (CA). In an elegant approach, NOVOMODO provides a scalable solution that requires only occasional contact with the CA [Micali 2002], and which has been adapted for use in mobile ad hoc networks [Schwingenschlögl et al. 2006] and grid computing [Sundaram and Chapman 2005]. Others have since proposed other solutions, also based on hash chains, with even better performance [Goyal 2007, e.g.,]. Because public keys may only be needed when introducing an SN to an MN, during which they can develop a shared secret used for future communications, this level of connectivity is likely sufficient. Some recent methods, such as Solworth [2008], may be useful when the relying party is a well-known, well-connected server (such as the HRS), allowing "instant" revocation in the case of a lost or stolen node.

Any security solution must be easy for the Patient to use. It must be intuitively easy for the Patient to introduce a new SN to their personal-area network (thus, to their MN) and thence to the HRS. The ubiquitous-computing community has developed several clever techniques for intentional pairing with secure key exchange [Kumar et al. 2009]. In one approach, the user literally shakes the two devices together; the devices communicate by radio but exchange key material derived from internal accelerometers [Mayrhofer and Gellersen 2007; Bichler et al. 2007]. In another approach, two nearby devices obtain key material from simultaneous observations of their RF environment [Varshavsky et al. 2007a, 2007b]. Another RF-based approach uses physical properties of the wireless channel between two nodes to derive key material [Mathur et al. 2008; Jana et al. 2009]; since the wireless channel fades quickly with distance, they show that an eavesdropping adversary cannot obtain the key, nor can an active adversary spoof one of the parties. Another approach uses physiological data as key material [Cherukuri et al. 2003].

Finally, since mobile devices like the MN and SN can be easily lost or stolen, secure key storage is an important challenge lest the keys become available to an adversary. We delve into this issue below in Section 5.1.3.

Anonymity. In most settings, such as treatment, billing, and health management, the Patient is clearly identifiable to anyone using the data. As we store more health data

electronically, there are huge opportunities for broad secondary uses: most usefully for research and for quality improvement. In that context, the Patient data can and should have control over whether their data may be used for secondary purposes (privacy properties P3, P6, P7), and their data must be anonymous (protecting properties P7, P9).

The HIPAA Privacy Rule states that covered entities may use or disclose PHI that is de-identified without restriction [HIPAA 2010]. Covered entities that seek to release such PHI must determine that the information has been de-identified using either statistical methods to verify de-identification or by removing certain parts of the PHI as specified in the Rule. Under the Rule, a covered entity can de-identify PHI by removing all 18 elements that could be used to identify the Patient or the Patient's relatives, employers, or household members. The Rule also requires the covered entity to have no actual knowledge that the remaining information could be used alone or in combination with other information to identify the Patient. Some of the 18 identifiable elements are the Patient's name, geographical information such as ZIP code, phone number, all elements of dates except the year, and biometrics. The HITECH Privacy Rule does not add anything new to this section of the HIPAA Privacy Rule.

It is notable that the ZIP code and birth date are two of the identifiable elements that covered entities must remove from PHI to de-identify it; Sweeney showed that it is possible to re-identify an individual by combining health data that includes ZIP, gender and birth date with a voter list that also contains these three elements [Sweeney 2002]. Sweeney describes a model called *k*-anonymity and policies to deploy the model such that an individual's data is indistinguishable from that of at least $k-1$ other individuals. Many others have built on this idea, with noted extensions such as ℓ-diversity [Machanavajjhala et al. 2006].

Health-data anonymization topic is an active area of research, with extensive literature; Appari and Johnson [2010] provide a useful overview. They cite several research efforts aimed at anonymizing Patient data: global and local recoding [Samarati 2001], Confidential Audits of Medical Record Access (CAMRA) [Malin and Airoldi 2007], microaggregation [Domingo-Ferrer et al. 2006] and data perturbation [Muralidhar and Sarathy 2005]. They also cite an interesting research effort by Riedl et al. [2007], which is to create a "secure pseudonymous linkage between Patient and her health record that will allow authorization to approved individuals, including healthcare providers, relatives, and researchers". Molina et al. [2009] explore the use of homomorphic encryption techniques to enable computation of aggregate functions over encrypted patient records.

Location Privacy. In some mHealth settings, the sensor data will be tagged with the current location; for example, a Patient desiring to be more physically active may collect data about their walking, running, and bicycling activity. The Patient may benefit from a detailed record, showing routes, maps, and dates [Eisenman et al. 2009], but may be uncomfortable sharing such detailed location records with their physician or other health services [Sankar and Jones 2005]. Appropriate use of anonymization is essential when dealing with location records (related to property P9). There are numerous studies demonstrating the risk of de-anonymization (that is, re-identification) of users even in "anonymized" location traces [De Mulder et al. 2008]. Location anonymity is particularly challenging if the adversary has external knowledge—perhaps obtained by observing or tailing the user—to combine with the anonymous location traces. Golle and Partridge [2009] show how easy it is to re-identify users in an anonymized location trace by recognizing home/work location pairs, even when locations are specified only by census block or census tract.

Location privacy is an active area of research; a thorough survey is beyond the scope of this article. There are many techniques to retain anonymity while releasing

information about your location, or about a sequence of locations (a path), each reflecting different use cases and different privacy goals; we discuss a few examples here. Some consider anonymous use of location-based services, for example, using techniques for blurring location and time [Mokbel et al. 2006; Gruteser and Grunwald 2003]; these methods have been extended to support personalized *k*-anonymity [Gedik and Liu 2008]. Hoh and Gruteser examine "path confusion," where the goal is to release a path (a series of location points) and yet retain anonymity; they do it by fuzzing points so that two users' paths sometimes "cross" [Hoh and Gruteser 2005; Hoh et al. 2007]. One limitation of these approaches, however, is that some require the user to trust a location service with the implementation of these methods, and thus with the responsibility for protecting the user's privacy. Sun et al. [2009] provides a framework that supports four methods to provide location information while maintaining privacy.

Even in an application that does not require the disclosure of location data, there is always the concern that wireless mobile devices might be tracked by the network infrastructure. Pang et al. [2007] recently showed that user identity might be recovered from 802.11 network traffic even if the MAC and IP addresses are changed periodically, so they later propose an anonymity-preserving variant to the 802.11 MAC protocol called SlyFi [Greenstein et al. 2008]. As we discuss below, the situation gets dire when an adversary uses fingerprinting techniques. Strong anonymity of network identity is still a difficult challenge.

5.1.2. Access Threats. In the next section of Table I we explore threats related to unauthorized access to PHI, whether in the MN or the PHR. Inappropriate access can lead to violations of privacy properties P3, P6, P7, P8, and P9. With good accountability mechanisms (P10), it may be possible to detect inappropriate access. The first threat comes from the Patient himself or herself, because (under the definition of health information privacy) the Patient has a right to control the collection, use, and disclosure of PHI (P3); if the Patient fails to express their consent in a way consistent with their actual preference, for whatever reason (perhaps because of a poor interface, P5), they may allow broader-than-intended collection, access or disclosure (P6, P7).

The Patient may also obtain more access than allowed, due to technical failures. Insiders may "peek" at Patient data, out of curiosity, or with the intent to harm the Patient (e.g., an employer who snoops on employer-provided PHR and fires workers with expensive conditions) [Appari and Johnson 2010; Sinclair and Smith 2008]. Outsiders may break into Patient records, which may lead to embarrassment (e.g., exposing a Patient's psychiatric data to his divorced spouse) [Messmer 2008].

Several of these threats involve the modification of health records. In a PHR, Patients (or insiders [Sinclair and Smith 2008]) may mistakenly modify their data if the access-control policies are too permissive, or if the mechanisms too easily allow mistakes. Insiders may modify PHI intentionally, to obtain reimbursement via insurance fraud [Dixon 2006]. Outsiders may also modify a Patient's PHI, for fraud or malice [Messmer 2008].

In this section, we survey work on consent management (allowing the Patient to determine what access is permitted), access control (policies and mechanisms), auditing (to support detection of violations), and data integrity.

Consent Management. A common legal requirement is that PHR service or health-care providers obtain consent from a Patient before disseminating her medical records to other providers or to entities such as a marketing department, medical researchers, or public-health officials. Informed consent is the most common way to think about properties P1–P3. The Markle Foundation's Connecting for Health initiative provides a list of considerations for obtaining consent [MFC 2009]. It categorizes consent into General Consent (an umbrella of privacy and terms of use policies) and Independent

Consent (indicating specifics separately from general agreement). Presenting the privacy policy in an understandable format—such that the Patient can set her preferences to provide or withhold consent—is a major challenge. Bellman et al. [2001] describe risks to Patient privacy due to the manner in which "opt-in" or "opt-out" questions are posed on a consent form. To address some of these issues, the US Department of Human and Health Services (HHS) proposed a format to describe a healthcare provider's privacy policy, based on the well-known "facts-at-a-glance" table seen on food products [HHS 2009] (an approach later shown to be effective [Kelley et al. 2010]). Halamka proposes the use of three flavors of consent: opt-in, opt-out and quilted (a subset of Patient data is shared only if the Patient explicitly consents, the rest is shared by default) [Halamka 2008]. The Healthcare IT Standards Panel (HITSP), a public-private partnership between US governments (federal and state) and health-related organizations, has released an interoperability standard (IS 03 Consumer Empowerment) specifying the management of machine-interpretable "consent directives" that Patients may issue to healthcare organizations on the access, collection, use and disclosure of their data [HITSP 2008]. It remains a challenge, however, to identify a clean and usable interface (P5) for consent management [Prasad and Kotz 2010].

Not all agree that consent is the right model for protecting Patient privacy; Nissenbaum [2004] proposes *contextual integrity* as a different way to look at the challenge of healthcare privacy.

Access Control. A mechanism for controlled access to a Patient's PHI, which restricts access to only legitimate entities, is necessary to ensure Patient privacy (affecting properties P4, P6, P7, P8). Standards bodies in the US, such as HL7 Standards Development Organization, have chosen the Role Based Access Control (RBAC) model [Ferraiolo and Kuhn 1992; Sandhu et al. 1996] to enforce access control in traditional healthcare IT systems. The RBAC model fits well in an organized healthcare setting, such as a hospital, because each entity in a hospital has a specific role and follows a well-defined hierarchy. However, certain features, such as "break-the-glass" to override access control rules in medical emergencies, do not exist in traditional RBAC [Motta and Furuie 2003; Covington et al. 2000]. Furthermore, RBAC is difficult to manage. Identifying roles and managing role membership is difficult in large organizations; some new "role-mining" methods attempt to help [Frank et al. 2009]. Finally, RBAC is not "privacy-aware"; access is either granted or denied. Ni et al. discuss how to extend standard RBAC to make it "privacy-aware" and enforce authorizations at a finer level of granularity [Ni et al. 2007a, 2007b]. Thus, privacy-aware RBAC authorizes a subject (e.g., a healthcare provider) to access an object (e.g., a Patient's record) only for a specific purpose (e.g., diagnosis) if it meets specific conditions (e.g., Patient's consent obtained) and obligations (e.g., notifying the Patient).

An important question related to the mHealth setting regards the "role" of the Patient. The Patient should not be able to delete or modify some types of entries in her PHR/EHR (e.g., readings from ECG sensors), but should be able to modify some other (e.g., diet information). It is not clear whether an RBAC model can enforce such fine-grained access control.

In protecting confidentiality of certain types of PHI via encryption, there are two issues to address: (i) who encrypts and decrypts the data: the Patient's MN, the PHR system, or some third party? (ii) how are encryption keys managed for entities that have access to Patient data, such that even if two independent entities collude they cannot obtain any new PHR data? Sahai and Waters [2005] and Goyal et al. [2006] describe *attribute-based encryption* for fine-grained access control, which associates data to be protected with an access-control policy and encrypts them with a public key personalized to each recipient of the data; the recipient must possess a secret key

derived from a set of attributes that satisfy the policy to decrypt the data. Sahai and Waters [2005] coined the term "functional encryption" to refer to this general body of work, which seems applicable to the problem of fine-grain PHR protection. However, the problem of key management and distribution remains and must be addressed, for example, by the use of a key authority. One potential simplification is the use of identity names as keys, with *identity-based encryption* (IBE) [Mont et al. 2003; Tan et al. 2009; Martin 2008]; Sahai and Waters [2005] allows a biometric reading to be used as the key.

Another aspect of access control, important in the mHealth setting, is the capability to securely access PHI on PHR sites using mobile phones. One such example is AllOne Mobile [AOM 2009], an application that can allow consumers to securely access their PHR data through their cell phones and PDAs. AllOne Mobile caches the Patient's PHR in a secure environment on the phone and synchronizes it with their online PHR. Patients can manage and share their PHRs with physicians and hospitals quickly and easily. AllOne Mobile wirelessly downloads a small application to the mobile device, which accesses an existing PHR stored in a remote server. The PHR can also be updated by the Patient using the mobile device. Using industry standards such as the Continuity of Care Record (CCR) and Health Level Seven (HL7), AllOne Health will also connect with Microsoft HealthVault, allowing Patients to access their PHR through their mobile phones.

Similarly, Anvita Health [AH 2009] has embraced the Android phone platform by rolling out a mobile interface for Google Health. With the new application, Patients and their physicians can have real-time access to electronic medical records stored on Google Health.

Auditing. Both HIPAA [Scholl et al. 2008] and HITECH [HITECH1 2009] require users within the healthcare provider's organization to be held accountable for their actions when handling Patients' protected health information (property P10 and supporting P11). The newer HITECH act gives covered entities a choice to maintain accounting of electronic disclosures by their business associates or to provide a list of business associates, who would then provide accounting details to individuals [HITECH1 2009; HITECH2 2009]. There are different approaches to maintaining audit controls for such information; one approach [IHE 2009] specifies a profile for the Audit Trail that contains sufficient information to answer questions such as: "For some user: which Patient's PHI was accessed? For some Patient PHI: which users accessed it? What user authentication failures were reported?" Such approaches could help administrators mitigate insider threats by ensuring detection of unauthorized access and illegal disclosure of PHI; they could also help detect attempts by hackers to break into a PHR system and help administrators detect potential vulnerabilities in the system, for example, ensuring the system fails unsafely when a specific sequence of transactions is executed. In the mHealth context, it may be necessary to implement audit mechanisms on the MN.

Data Integrity. The HIPAA Security Rule (Section 164.312(c)(1) Integrity) states that covered entities must "implement policies and procedures to protect electronic PHI from improper alteration or destruction" [Scholl et al. 2008] (important for property P8). In non-medical settings, applications usually verify data integrity, either by means of a checksum or a hash, before using the data. If the integrity check fails, the application reports an error and terminates without processing the data. In the healthcare setting, databases storing Patient data must implement integrity computation and verification functionality. Blough et al. [2008] have proposed the use of Merkle trees to store a PHR.[3]

---

[3]Merkle trees [Merkle 1982] are a type of data structure, useful to store files along with their integrity values. A Merkle tree is an *m*-ary tree in which all leaves contain the actual data that are related (e.g., parts of a file, files in a directory) and each nonleaf node contains a hash of the subtree of which it is a root.

Each type of data in a PHR, such as lab results, MRI images, or diet information, is stored as a separate leaf in the tree, with related data forming a subtree. They claim that Merkle-tree based PHRs enable Patients to selectively disclose data to providers by choosing the appropriate sub-tree for disclosure and also for providers to verify that their data were not tampered within the PHR.

5.1.3. Disclosure Threats. In the final section of Table I we explore threats related to the disclosure of PHI, including data at rest and data in transit (affecting properties P7, P9). PHI stored in a PHR may be disclosed inadvertently, by the Patient or other insider, if malware or file-sharing software on their computer makes it available [Johnson 2009]. Similarly, Patients or other insiders who share access credentials may allow data disclosure beyond what was intended by the act of sharing. Finally, insiders (or outsiders) with access to PHI may disclose that data to others for financial gain or to embarrass the Patient [Appari and Johnson 2010]. These threats affect property P7, and indirectly property P3.

In an mHealth system, much interesting PHI is created by sensors, and stored at least temporarily in the SN or MN; these devices, if lost or stolen, could lead to disclosure of the sensor data. Worse, such devices contain cryptographic keys that may allow one to decrypt current, prior, or future communications by that device. Examination of an MN may also allow one to determine what sensors are in use by the Patient, and what sensor data is to be collected (property P9). A Patient may not want an outsider to know, for example, that she is wearing an ECG sensor.

Furthermore, an mHealth system is inherently distributed, and network communication among the entities may be captured and analyzed by an outsider. Even if the traffic is encrypted, device types may be discernible (through network fingerprinting methods or address formats), and analysis of network traffic patterns may identify the sensors or sensing tasks (property P9).

In this section, we survey work related to secure data transmission, device presence, and device compromise and theft.

Secure Transmission. The HIPAA Security Rule (Section 164.312(e) Transmission Security) states that covered entities must "implement technical security measures to guard against unauthorized access to electronic protected health information . . . transmitted over an electronic communications network" [Scholl et al. 2008]. The 2009 HITECH Act extends this rule to business associates [HITECH1 2009; HITECH2 2009]. Although HIPAA's rule covers communication between HIPAA-covered entities, our concern here is an adversary who wishes to obtain confidential medical information from observing the network communications between the MN and its SNs, or between the MN and the distant health services. In the mHealth setting, we must assume the use of wireless networks and open standards. There are five fundamental challenges here, all affecting properties P7 or P9.

First, the adversary may inspect the wireless-network packets and obtain sensitive medical data; this problem can be resolved by encrypting all communications with a secure encryption method and an appropriately strong encryption key. Most emerging services use HTTP over SSL, but we know of one approach leveraging SIM-card support in mobile phones [Weerasinghe et al. 2007]. Key management remains a challenge, however, as noted above.

Second, even if the wireless-network traffic is encrypted, in some settings it is possible for a clever adversary to use traffic analysis—the study of the size and timing of network packets—to determine characteristics of the traffic [Wright et al. 2010]. It may be possible, for example, to determine the type of sensor node from the pattern of its communications, or the type of medical application by observing the pattern of MN communications [Srinivasan et al. 2008]. Although there is literature on mechanisms

to defeat traffic analysis [Wang et al. 2008; Wright et al. 2009] (for example), we are unaware of anything specific to the mHealth context. Furthermore, many solutions require the addition of delays (to defeat timing analysis) or padding (to defeat packet-size analysis) [Srinivasan et al. 2008], which may harm system performance and increase resource usage.

Third, the adversary may use physical-layer or link-layer fingerprinting methods to identify the device type. In general, fingerprinting techniques observe subtle differences in the network behavior of devices, because of implementation differences across manufacturers, models, revisions, or even individual devices. Research has shown that it is possible to distinguish individual Wi-Fi network cards based on observations of their physical-layer (radio) behavior [Brik et al. 2008], or their link-layer behavior using passive [Franklin et al. 2006] or active [Bratus et al. 2008] methods. It is extremely difficult to combat these methods, particularly at the physical layer.

Fourth, because the wireless medium is open, an active adversary may inject frames or may selectively interfere with (cause collisions with) wireless frames. These methods may enable the adversary to create a man-in-the-middle situation [Dai Zovi and Macaulay 2005], to use link-layer fingerprinting methods [Bratus et al. 2008], or to compromise the devices in a way that divulges their secrets. Indeed, there are increasing concerns (and demonstrated attacks) regarding the wireless communications of implanted medical devices [Halperin et al. 2008a, 2008b].

Fifth, the devices must be somehow introduced to each other. The Patient's constellation of devices includes personal devices (mobile node, wearable sensor nodes), family devices (such as a bathroom scale), and clinical devices (when visiting the doctor). It must be easy for the Patient to establish a security relationship with these devices, without confusing his devices with those of his housemate [Cornelius and Kotz 2011]. Several groups have proposed possible protocols [Andersen 2009; Baldus et al. 2004; Malasri and Wang 2008], for various settings; some use the body itself as a communication medium [Barth et al. 2008; Garcia-Morchon et al. 2009; Garcia-Morchon and Baldus 2008]. Another proposes the use of identity-based encryption [Tan et al. 2009], with an implementation on mote-class sensor devices.

Device Presence. A Patient may consider the fact that she is using personal medical sensors to be private information (property P9); she may not want an employer to know, for example, that she is wearing a fetal monitor. The challenge, then, is to allow MN-SN communication, without exposing to an eavesdropper the fact that they are *medical* devices let alone which types of medical devices. The challenge is especially difficult given that many mHealth scenarios expect the MN to discover (or re-discover) SNs that are not currently with the Patient, but which need to be contacted when they become in range (such as when a Patient arrives home and re-connects with their blood-pressure device).

Most of the relevant work relates to network-identifier privacy, including Wi-Fi [Pang et al. 2007], Bluetooth [Singelée and Preneel 2006; Wong and Stajano 2005], and presumably Zigbee. Mare et al. [2011] proposes a new privacy-sensitive protocol for low-power body-area sensing in mHealth scenarios; although not aligned with any standard, it provides strong privacy properties. It remains to be seen whether there is a standards-compliant solution in which the link-layer identifiers (and other fields related to link-layer discovery) can be constructed to not leak information about sensor type. As shown in the fingerprinting discussion, above, even that may not be sufficient to solve the problem.

In a wireless personal-area network, like that between the MN and its SNs, there is a similar risk that the device "discovery" messages may allow an attacker to determine sensor types—even in standard protocols like Bluetooth [Mare and Kotz 2010].

A technique for Bluetooth identifier privacy may be helpful here, because it masks the identity of the nodes that are communicating [Singelée and Preneel 2006]; if such a protocol were used for session establishment, developing a shared secret key, then other sensitive metadata can be shared over a secure channel.

Ultimately, this particular challenge will be difficult to solve. The mere presence of a wireless signal emitting from your body does not necessarily tell an adversary that you have a medical device, of course, because of the increasing prevalence of personal wireless devices; the presence of a wireless device seeking to "cloak" its identity, however, may raise suspicions.

Device Compromise, Theft. The Patient's MN may be compromised, for example, by an email-borne virus. The MN or SN devices may be lost or stolen. In any case there are several risks. The adversary may obtain access to personal health information, or learn the type of sensor nodes, both of which may expose information about the Patient's medical condition (affecting properties P7, P9). Moreover, any key material may be obtained by the adversary, potentially allowing the adversary to decrypt previous, current, or future communications between the Patient's MN and SNs, or between the MN and the health records system (HRS). Furthermore, the key material may enable the adversary to inject false data into the MN or health records system, or even to reconfigure the Patient's other devices (affecting P8). Finally, the key material may enable the adversary to decrypt data stored in the health records system (affecting P7). The specific risks depend on the protocols used for data transfer, on the encryption methods used, and on the security of key generation, distribution, revocation, and storage.

We have not seen any complete solution to this problem in the literature, but any solution requires solving two fundamental problems: key management and secure storage on mobile devices.

Key management is discussed above, but in the context of device loss there are two important considerations. First, the keys used by the smallest devices—most likely to be lost—should have limited lifetimes, limiting the damage caused by key compromise. Similarly, the keys with the most importance—authenticating SNs to the MN, authenticating MNs to the HRS, or the Patient to the system, should expire or be easily revocable. A short key lifetime, however, requires the generation (and distribution and validation) of new keys frequently, which may be difficult in loosely connected networks or in low-power devices. Key revocation, on the other hand, requires good connectivity and the storage of revoked-key lists (see previous section). Again, some research proposes key-management approaches given these constraints [Malasri and Wang 2008].

Persistent storage in the MN or SNs may contain sensor data, configuration information, and cryptographic keys. An adversary with physical access to either node may be able to extract this information, possibly through an internal memory bus or connector. Such attacks have been demonstrated on common "mote" sensor nodes [Becher et al. 2006]. The technical challenge is to make such information access difficult; although software-based solutions (such as hypervisors) provide some protection [Gilbert et al. 2010; Kleidermacher 2008], a strong solution rests squarely on methods for tamper-resistant trusted hardware. Since it may be expensive to build a large, tamper-resistant persistent memory, most approaches use conventional storage technology (flash or disk) and encrypt the sensitive data stored there; then the problem reduces to secure storage of the encryption key. Although a Trusted Platform Module [TPM 2009] (TPM) is common today in many of today's laptops, a standard is still under development for mobile-phone class devices [MPWG 2009; MTM 2008]. Even that solution will likely be too expensive for SNs. One paper has demonstrated the addition of TPM hardware to a mote-class device [Hu et al. 2009], but the authors have not yet demonstrated its use for secure storage.

A possibly more insidious threat is the unintended disclosure of sensor data by applications installed by the user, including applications unrelated to healthcare. The security community is just beginning to address this threat [Enck et al. 2009].

5.2. Human Issues

When using mHealth systems, Patients may be collecting large amounts of information about their health and their activities, throughout daily life, and we expect they will want control over what is shared with their health provider or others. It can be difficult for Patients to decide what information to share, and these choices can have important health consequences [Sankar and Jones 2005]. The right interface, in each setting, is important.

In their seminal paper "The Protection of Information in Computer Systems", Saltzer and Schroeder [1975] state the principle of *psychological acceptability* as follows: "It is essential that the *human interface* be designed for ease of use, so that users *routinely* and *automatically* apply the protection mechanisms *correctly*. Also, to the extent that the *user's mental image* of his *protection goals* matches the mechanisms he must use, mistakes will be minimized. If he must translate his image of his protection needs into a radically different specification language, he will make errors" (emphasis added).

In this section, we first survey research that discusses the challenges and opportunities of creating human interfaces to enable privacy management, which is privacy property P5. Then, we survey studies and research efforts that attempt to understand users' perception of privacy and its importance to them.

5.2.1. Human Interfaces for Privacy Management. Recalling the NCVHS definition of health information privacy, any entity that manages PHI must provide the Patient with interfaces that enable her to exercise her right to control the collection, use, and disclosure of her PHI, after evaluating the request based on relevant information, for example, who is the recipient or what is the purpose. That is, the Patient must exercise informed consent.

Friedman et al. [2005] describe a "conceptual model of informed consent for information systems". Their model consists of five components: disclosure, comprehension, voluntariness, competence and agreement. The first two components contribute to the user becoming informed, while the remaining three are criteria to determine the user's consent.

Informing the user can be remarkably challenging, because users tend to overlook or misunderstand security and privacy notices. Cranor provides a framework for thinking about how to present security warnings to the user so that they can effectively take notice, understand the warning, evaluate options, and take action [Cranor 2008]; this framework may be useful in designing mechanisms for informed consent. Some creative proposals have been developed, including the use of a "nutrition label" format for privacy policies [Kelley et al. 2010] and the concept of "privacy icons" in which websites would display their privacy practices as a set of intuitive icons (sample shown in Figure 2).

Karat et al. [2005] identify key challenges in designing privacy and security systems that are *usable*. Accordingly, the first key issue is to recognize that when using a system, e.g., to update their PHI or order medications, a user's main goal is not to protect their privacy but to complete the task at hand; thus, privacy controls must be as transparent as possible to the user, but also accessible and usable whenever the user needs to understand what is happening or needs greater control over the situation. Second, the design should make the privacy system accessible to all types of users, not just technical experts. Third, designers must recognize that users will not use a privacy system that is too complex to understand, putting them at greater risk of exposure

Fig. 2. Sample privacy icons (reproduced with permission by Aza Raskin, http://www.azarask.in/blog/post/privacy-icons/).

than if they used a simpler, less sophisticated system. Finally, the design should allow its operators to easily update it to meet changes in regulatory requirements.

Cranor discusses challenges of creating interfaces that enable users to express privacy preferences on the Web [Cranor 2003, 2005]. Based on results of a user survey (in the US) that provided information on "aspects of privacy policies that would likely be of most interest to users", Cranor [2005] determined that the three most important areas to users were: the type of data collected, how data would be used and whether data would be shared. Also, users were most concerned about financial and medical data collected, and the use of their data for telemarketing or marketing lists. The paper reports that focus-group users evaluating early prototypes of automated privacy preference agents had two seemingly contradictory requirements for the interfaces: the users "wanted the interface to be extremely simple, but they also were reluctant to have their choices reduced to several pre-configured settings such as high, medium and low".

Indeed, it may be possible to use machine-learning methods to discover users' privacy preferences, for example, by providing the user opportunity to provide feedback on the system's automated decisions [Kelley et al. 2008] or by allowing the user to inspect their location history and annotate how they would (or would not) have wanted that information shared [Ravichandran et al. 2009].

For an extensive survey of HCI (human-computer interfaces) and privacy, though not specifically about healthcare or mHealth, we refer readers to Iachello and Hong [2007].

5.2.2. Privacy Perception and Attitudes. Research and human studies have established that concern about one's privacy is not uniform among the human population, even in the case of "high risk" information such as e-commerce data [Ackerman and Mainwaring 2005]. Privacy concerns range from unauthorized entities observing one's data, to entities aggregating personal data to create a profile, to unauthorized entities (including those unrelated to the original recipients) actively modifying personal data. Ackerman and Mainwaring note that a persistent finding in human studies and research is that users in the US cannot be considered as a homogeneous group with regard to privacy concerns [Ackerman and Mainwaring 2005].

More recently, a group of mHealth researchers interviewed 24 users of a sensor-based fitness application [Klasnja et al. 2009]. They were asked about the data collected, and about hypothetical extensions to the system; the results show that their privacy concerns depended on what data was being collected, how long it was stored, the context in which it was collected (e.g., at work or at home), and the perceived benefits of the resulting application. No one privacy policy fits all settings, let alone all users.

Finally, the Patient's privacy wishes—and the practical options available to them— may depend on the nature of their medical condition [Prasad et al. 2011]. A Patient who is struggling with a debilitating disease may be willing to share more information with their physician, than would a healthy Patient who is seeking routine care. Conversely, a physician may (correctly) insist that the Patient share a large amount of detailed

information, to be able to properly diagnose and manage treatment of the Patient's condition. The challenge here will be to help Patients (and clinical staff) thoughtfully navigate the boundaries of what information can and cannot be withheld from the clinical team.

5.2.3. Summary. User perception, awareness and concern about privacy of their data, including medical information, vary widely across the world. While individualist societies such as those in the US place a high premium on sharing personal data with others, collectivist societies such as those in India do not [Kumaraguru and Cranor 2006]; further, western societies and developed regions of the world have a greater amount of experience with the Internet and its problems with regard to privacy than developing nations, which results in differing views of privacy. Therefore, there is no "one size fits all" solution to the problem of privacy management. Even within a geographic region, such as the US, the problem of creating a usable interface that enables people to express their privacy preferences remains open. This challenge results from the many different types, usages, and recipients of data, and the operations that can be performed on them. In the context of privacy management for PHI, there is a need to bound the problem to a specific subset of types, usages, recipients, operations and perceptions to make it more tractable.

6. OPEN RESEARCH QUESTIONS

Given the motivation provided in Section 1, the privacy properties in the list of properties from Section 3.1, and the prior work surveyed in Section 5, in this section we lay out a series of specific research questions focused on privacy in mHealth systems.

Although mHealth raises many challenging security and privacy research issues, in this survey we focus on the privacy-related challenges and on patient-centric mHealth technology. Thus, while we recognize the importance of security and data integrity, for example, we address these issues only to the extent that their solutions have an impact on privacy. We identified eight major research topics, each inspired by the desired privacy properties (Section 3.1) or by the privacy threats (Table I).

(1) *Consent Management.* How can the Patient use their MN to easily manage consent, that is, express preferences over collection, dissemination and retention of PHI, and make consent decisions when requests occur?

(a) What kinds of consent do Patients need to express, and when? At what level of detail do they wish to control their privacy? If Patients are given fine-grained control, it provides flexibility with a heavy cognitive burden, and it may adversely impact the usefulness of the data to a healthcare provider who must interpret a medical record full of holes. On the other hand, coarse-grained choices may be too blunt to give the Patient real choice: all too often, today's systems provide an all-or-nothing consent form, which the Patient must sign in order to obtain any service.

(b) What interface effectively informs the Patient what information will be collected (and why), where the data will reside, what will be disclosed and to whom, for what purpose it will be used, and for how long it will be retained? Keep in mind that some data may be kept for decades. Furthermore, how can we provide the Patient an understandable assessment about both the privacy risk and the potential benefits of information sharing, to enable him to make informed consent decisions?

(c) What interface allows the Patient to easily express preferences about what data may be collected (and when), what may be disclosed and to whom, for what purpose it may be used, and for how long it may be retained?

(d) Can machine learning algorithms help to define, or refine, a Patient's privacy preferences?

(e) Can we automate consent management? If so, how should requests for PHI be encoded, how should privacy preferences be encoded, and what algorithms allow software to automate consent decisions? How can the *context* of the Patient, the technology, and the Consumer help to support automated consent management? For example, prompt the user to update privacy preferences when adding a new SN, simplify sharing of PHI while meeting with a physician, and automate consent in an emergency scenario.

(f) What mechanisms can enable the Patient (really MN on behalf of the Patient) to authenticate the person or organization requesting the PHI?

(g) Regional Variations: Identify regional variations in law, economics, and culture in terms of privacy, healthcare and mobile technology, individually and in combination. For example, in many cultures it is common for the extended family to assist a Patient at home, a setting in which many individuals may have (and need) access to the MN and perhaps to make privacy decisions on the Patient's behalf. We must recognize that many individuals may act in a "custodial" or "caregiver" capacity with respect to the Patient, and to the Patient's data. Furthermore, different legal and economic structures pose different incentives to providers, insurers, and clinical staff regarding Patient privacy. How do we design a system to easily accommodate these regional differences?

(h) Over a lifetime, a patient's medical records may be managed by her parents, then by herself, then (if she becomes infirm) by her spouse or children. How do these transitions occur, and relate to legal concepts such as durable healthcare power of attorney? May a patient block certain records (for e.g., about an attempted suicide or abortion) so that her children or spouse may never see them?

For all of these questions, consider the unique challenges of human-computer interface on a mobile platform. This interface must be extremely simple, intuitive, and unobtrusive. We seek to explore a broad range of interface modes, including web and social-network abstractions as well as audio, visual, and touch interfaces. Ultimately, this topic requires significantly more study to determine whether Patients need (or even want) rich privacy interfaces, and whether there are differences between what they want in home-health settings and in hospital settings. Simpler models may be suitable for certain situations, such as health-care social-network applications. Any solution needs to recognize that the medical data collection in an mHealth setting is more pervasive than a traditional clinical setting and also includes collection of Patient context information. These differences add complexity to the solution in an mHealth setting.

(2) *MN Architecture*. Identify which of the privacy properties (Section 3.1) require support on the MN, and what privacy-related requirements they impose. How should MN hardware and software architecture change to help protect Patient privacy and enable them to manage privacy? Specifically, what hardware and software enhancements would help

(a) to enforce Patient's privacy preferences,

(b) to protect contents in the MN (PHI, keys and software),

(c) to preserve privacy of Patient context (location, device presence, communication, activity, etc.),

(d) to create a secure execution space on the MN for contextual health-sensing and PHI-handling applications,

(e) to allow multiple software and services to co-exist on the MN, without conflict, and to enable software updates to be pushed and securely installed, and

(f) to easily manage user authentication, data collection, and manageability (e.g., remote disable and remote updates).

In what way does *privacy* require enhancements that may be different than those required to support other security and functional properties?

In what way does PHI (or the healthcare setting) require enhancements different than those in other applications, such as e-commerce or entertainment?

Finally, any solution to these problems must consider the resource constraints (memory, CPU, bandwidth, energy, and human interface) available on MID-class devices.

(3) *Enforcing Control over Data*: How can we *enforce* control over PHI?

(a) How can we securely bind policies to PHI as it flows to different entities? Can we use DRM-like techniques on PHI? Does PHI (or the healthcare setting) require DRM solutions different than those in other applications?

(b) How can those policies enforce control over PHI data (for example, auto-destruct data after a specified time limit, limit the number of views, limit the number of copies, or identify the entity responsible for copying data)?

(c) How can we ensure that PHI is used only for the specified purpose and monitor compliance? In general, where do we enforce policy—at the MN, at the HRS, or both?

(d) How can we accomplish these types of control with an interface that is *usable* by clinical staff in clinical settings?

(4) *Data Identity*. Consider mechanisms to verify that the sensors are attached to the correct Patient. What solutions provide reliable Patient identity verification *and* preserve Patient privacy in the process? Furthermore, how can the sensor data be labeled with Patient identity in a way that allows data Consumers to verify the assertion that this data came from this Patient, while preserving the Patient's privacy? In other words, any solution to the problem of verifying Patient identity, labeling the data with that identity, and asserting validity of the label, must be designed with Patient privacy in mind.

(5) *Anonymization*. What are effective algorithms to anonymize PII before disclosing it to another party, for example, for research or for a medical opinion? Under what circumstances should anonymization be performed at the MN rather than in the HRS or other back-end servers? Consider that additional context information may make anonymization more difficult in an mHealth setting. Are there usable mechanisms to allow Patients or health providers to control what information is disclosed?

(6) *Accountability*. What mechanisms can be used to support accountability (i.e., make Consumers of PHI accountable to the Patient for using PHI according to her preferences) and nonrepudiation?

(a) How can we effectively inform the Patient of the risks of a privacy breach?

(b) How can we effectively create audit logs for all access, addition, deletion, and modification of PHI?

(c) What interfaces can enable the Patient to review audit logs in an easy and understandable manner? In particular, who accessed the data, and were they authorized to have access at that time?

(d) What interfaces allow the Patient to annotate or correct their sensor data or the inferences derived from that data?

(e) What mechanisms can help to remedy effects of security breaches or privacy violations?

(f) How can we effectively notify the Patients of a security or privacy breach, its consequences and remedial measures in an understandable manner?

(7) *Ecosystem*. What are the ecosystem support roles in an mHealth system? We initially identified five roles; policymakers, certification bodies, manufacturers, distribution, remote management, and a key-management infrastructure. What policy and legal frameworks need to be in place for them to serve these roles? What standards need to be developed, and what certification mechanisms can encourage and ensure compliance with standards?

(8) *Trade-offs*. Solutions to these problems involve many trade-offs, such as between anonymity and accountability, or Patient authenticity and privacy. An important research challenge, then, is to develop a conceptual working framework to help identify these trade-offs within a solution space, and choose an appropriate balance when making design decisions.

An extensive research agenda, indeed. There may be yet more challenges, but all of these questions must be addressed to develop an mHealth system that supports the principles in our mHealth framework (Section 3). Although some prior work, as discussed in Section 5, addresses aspects of each of the above research questions, none of the questions has been fully resolved in the context of mHealth. Furthermore, each of our privacy properties is addressed by one or more of the above research questions, and the properties are themselves linked back to the core principles we identified. We encourage the mHealth research community to internalize the principles, address these research questions, and build systems that support all the desired properties.

7. SUMMARY

In this article, we sought to identify the privacy challenges in mHealth systems, and survey the state of the art in the literature. We developed a conceptual privacy framework for mHealth, and use it to itemize the privacy properties needed in mHealth systems. In an extensive survey of the literature, we discuss the technologies that could support privacy-sensitive mHealth systems.

The technology of mHealth is an area of extremely active research and development, with increasing interest both from academic researchers and corporate product developers. We believe it is essential to consider privacy in the design and implementation of any mHealth system, given the sensitivity of the data collected. This article provides a thorough list of privacy properties, which for a developer are a checklist that should be considered in any design, and for a researcher generate questions about how to enable those properties more efficiently and more effectively. Interdisciplinary research is essential, since many of the research challenges involve both technology and human factors, or technology and policy issues. Finally, we note that the privacy concerns, and privacy-supporting technologies, should be considered in the context of the medical intent of the mHealth system.

The challenges we identify need to be addressed soon, because mHealth devices and systems are being deployed now. It will be far more difficult to retrofit privacy protections than to build them in from the start.

APPENDIX

A. PRIVACY FRAMEWORKS FOR HEALTHCARE

In the main document, we refer to several lists of privacy principles produced by various organizations. In this appendix, we quote those principles, adding numbers of ease of reference within this article. We include commentary to introduce and compare the frameworks. A version of this appendix appeared as a workshop paper [Kotz et al. 2009].

A.1. ONC National Framework (2008)

In December 2008, the Office of the National Coordinator for Health Information Technology (in the US Department of Health and Human Services) announced its *Nationwide Privacy and Security Framework for Electronic Exchange of Individually Identifiable Health Information* [ONC 2008]. This document, which we call the "ONC Framework", provides the newest and most authoritative privacy framework for healthcare released in the US, so we present it first. We quote their eight principles as follows, adding numbers for ease of reference within this article. Their document includes additional explanatory detail that we do not quote here.

ONC1. INDIVIDUAL ACCESS. Individuals should be provided with a simple and timely means to access and obtain their individually identifiable health information in a readable form and format.

ONC2. CORRECTION. Individuals should be provided with a timely means to dispute the accuracy or integrity of their individually identifiable health information, and to have erroneous information corrected or to have a dispute documented if their requests are denied.

ONC3. OPENNESS AND TRANSPARENCY. There should be openness and transparency about policies, procedures, and technologies that directly affect individuals and/or their individually identifiable health information.

ONC4. INDIVIDUAL CHOICE. Individuals should be provided a reasonable opportunity and capability to make informed decisions about the collection, use, and disclosure of their individually identifiable health information.

ONC5. COLLECTION, USE, AND DISCLOSURE LIMITATION. Individually identifiable health information should be collected, used, and/or disclosed only to the extent necessary to accomplish a specified purpose(s) and never to discriminate inappropriately.

ONC6. DATA QUALITY AND INTEGRITY. Persons and entities should take reasonable steps to ensure that individually identifiable health information is complete, accurate, and up-to-date to the extent necessary for the person's or entity's intended purposes and has not been altered or destroyed in an unauthorized manner.

ONC7. SAFEGUARDS. Individually identifiable health information should be protected with reasonable administrative, technical, and physical safeguards to ensure its confidentiality, integrity, and availability and to prevent unauthorized or inappropriate access, use, or disclosure.

ONC8. ACCOUNTABILITY. These principles should be implemented, and adherence assured, through appropriate monitoring and other means and methods should be in place to report and mitigate nonadherence and breaches.

We next consider some important predecessors and inspirations for the ONC Framework, in chronological order.

A.2. Health Privacy Project—Best Practices (2007)

The Health Privacy Project listed 10 "best practices" for employers who are developing personal health records (PHR) [HPP 2007]. (The Health Privacy Project has recently been adopted by the Center for Democracy & Technology, CDT.) We quote their list as follows, adding numbers for reference. Note that these principles refer to "employer" and "employee" rather than "provider" and "patient", because they are intended for employer-provided PHRs.

BP1. Transparency and notice. Employers should be transparent about their reasons for offering a PHR to employees and all policies that apply to the PHR. Employers should provide an Information Policy Statement or Notice that clearly lays out the

ways in which information in the PHR will be used and safeguarded. Employers should incorporate the Notice into their health benefit programs, and should make it available in a layered format—a short concise version to accompany a more detailed one. Employees should be informed of any updates to the policy.

BP2. Education. Employees should be educated about the benefits, functions, and content of the PHR. Information about the PHR should be communicated in numerous ways to build both knowledge and trust.

BP3. Employees can choose which content is included in the PHR. Employees should be able to determine the content of the PHR, including which providers and plans contribute to it. Employees should be able to annotate the records submitted by others, as well as to enter their own information, with employee-entered data marked as such. The identification of sources of all personal health information in the PHR should be readily apparent.

BP4. Employees control access to and use of the PHR. (A) Employees should control who is allowed to access their PHRs. Employers should not access or use employees' individually identifiable health information from the PHR. (B) Employees should choose, without condition, whether to grant access to personal health information within their PHRs for any "secondary uses". An audit trail that shows who has accessed the PHR should be easily available to employees.

BP5. Employees can designate proxies to act on their behalf. Employees should determine who, including family members and caregivers, should have direct access to their PHRs on their behalf. Where possible, employees should be able to grant proxy access to full or partial information in their PHRs, including access in emergency circumstances. Employees should also have the ability to revoke access privileges.

BP6. "Chain of trust": Information policies extend to business partners. The information policies and practices of employer-sponsored PHRs should follow the data through chain of trust agreements that require business partners to adhere to the employer's applicable policies and practices.

BP7. Data security. Employers should provide a strong level of security to safeguard the information in the PHR systems. A robust authentication process for access to PHRs should be required, in addition to an audit trail that shows who has accessed information and when.

BP8. Data management. Employers should ensure that the PHR systems they provide have comprehensive data management strategies that protect the integrity of the data and include data retention policies.

BP9. Enforcement and remedies. Employers should establish oversight and accountability mechanisms for adhering to their PHR policies and practices. Employers should put into place a mechanism to promptly notify employees of any inappropriate access to or use of information contained in an employee's PHR, identify the steps which have been taken to address the inappropriate activity, and make resources available to employees to assist them in addressing the effects of the inappropriate activity.

BP10. Portability. Employers should offer PHRs that are portable, to the extent feasible, allowing employees to maintain or move the PHR and/or the data it contains even after employment or coverage ends or changes.

The HPP Best Practices list contains aspects specific to PHR systems, such as Patient-entered data (BP3) and portability (BP10). There are some aspects specific to employer-provided systems, such as Education (BP2). BP5 mentions the concept of a "proxy", which is not mentioned by any of the other frameworks, except ONC4 (in the details). The "chain of trust" concept (BP6) is more explicit than in any of the privacy frameworks; this transitive application of privacy constraints does not explicitly appear

in the ONC principles, but we believe it is an essential feature. Indeed, we anticipate that there may be interesting technological methods for wrapping PHI with privacy policies before sharing with a third party. There is explicit mention of the requirement to notify the Patient of any inappropriate disclosure (BP9); the ONC Framework only mentions such notice in its detailed comments, and only as an example of a reaction and not as a requirement. The Common Framework mentions a similar requirement in its detailed comments about CF7.

Eight years earlier, in 1999, the Health Privacy Project released a set of "best principles" for health privacy [HPP 1999]. The document notes that their "principles are intended to establish a comprehensive framework". Many of the themes behind these principles resonate in the ONC Framework, but we want to call particular attention to HPP2, which notes that "All recipients of health information should be bound by all the protections and limitations attached to the data at the initial point of collection". This transitive application of privacy constraints does not explicitly appear in the ONC principles, but we believe it is an essential feature and deserves its top-level appearance in the HPP Framework. (This principle shows up again as BP6.)

A.3. Markle Foundation's "Common Framework" (2008)

The Markle Foundation launched a project "Connecting for Health", which brought together a wide range of stakeholders in developing a "Common Framework", a model for healthcare information exchange [MF 2008]. The Connecting for Health Common Framework (CF) describes both policy and technical principles for healthcare information exchange, and provides concrete prototypes of each: everything from patient consent forms to data-interchange formats and information architecture. The Center for Democracy & Technology later endorsed the policy aspects of the Common Framework in their own policy document [CDT 2008]. We quote the top-level description of the CF principles here.

CF1. Openness and transparency: Consumers should be able to know what information has been collected about them, the purpose of its use, who can access and use it, and where it resides. They should also be informed about how they may obtain access to information collected about them and how they may control who has access to it.

CF2. Purpose specification: The purposes for which personal data are collected should be specified at the time of collection, and the subsequent use should be limited to those purposes, or others that are specified on each occasion of change of purpose.

CF3. Collection limitation and data minimization: Personal health information should only be collected for specified purposes and should be obtained by lawful and fair means. The collection and storage of personal health data should be limited to that information necessary to carry out the specified purpose. Where possible, consumers should have the knowledge of or provide consent for collection of their personal health information.

CF4. Use limitation: Personal data should not be disclosed, made available, or otherwise used for purposes other than those specified.

CF5. Individual participation and control: Consumers should be able to control access to their personal information. They should know who is storing what information on them, and how that information is being used. They should also be able to review the way their information is being used or stored.

CF6. Data quality and integrity: All personal data collected should be relevant to the purposes for which they are to be used and should be accurate, complete, and up-to-date.

CF7. Security safeguards and controls: Reasonable safeguards should protect personal data against such risks as loss or unauthorized access, use, destruction, modification, or disclosure.

CF8. Accountability and oversight: Entities in control of personal health information must be held accountable for implementing these principles.

CF9. Remedies: Remedies must exist to address security breaches or privacy violations.

The ONC Framework covers all these principles. CF1 is covered by ONC3 (Openness and transparency), and ONC1 (Individual access). CF2 is covered by ONC5 (Collection, use, and disclosure notification) and ONC1. CF3 is covered by ONC5 and ONC4 (Individual choice). CF4 is covered by ONC5. CF5 is covered by ONC1 and ONC4. CF6 is covered by ONC6 (Data quality and integrity). CF7 is covered by ONC7 (Safeguards). CF8 is covered by ONC8 (Accountability). CF9 is covered by ONC8 and ONC2 (Correction). We believe that the Common Framework, though, is a crisper statement of the important principles; it is also more explicit in emphasizing that data should be used only for the purpose for which it was collected, and that Patients should have control over what is collected and to whom the data may be disclosed.

Furthermore, we see little in the ONC principles that is not covered in the Common Framework, except that ONC1 provides more explicit statement that Patients should have an easy method to obtain their PHI, ONC2 provides an explicit statement that Patients should be able to correct mistakes in their records, ONC3 explicitly states openness about policies and technologies, and ONC4 emphasizes that Patients should be able to make informed choices.

A.4. CCHIT's Certification Criteria (2008)

The Certification Commission for Healthcare Information Technology (CCHIT) is a non-profit organization that certifies healthcare information systems, and they recently released their certification criteria for PHRs. Although it appears that they have a process in place to determine these criteria, and that the process may enable and encourage updates to the criteria, we quote the current list of criteria as published in a booklet copyrighted in 2008 [CCHIT 2008].

CCHIT1. Consent. You should be in control of your personal health information and how it is used. PHRs that meet certification requirements must include safeguards that require you to give your explicit consent before your account is opened, or allow you to opt out of the service. It also must allow you to decide if your data can be collected, displayed, accessed, stored, released or disclosed.

CCHIT2. Controlling Access to your Information. Your PHR should give you the ability to decide what information is private and to restrict access to it. Your PHR provider must get your permission to gather or disseminate any information about you. You also decide who else can view information in your PHR, and limit the types of information that can be viewed.

CCHIT3. Conditions of Use. The conditions for using your PHR should be explicitly explained to you, and you have the right to challenge your PHR provider if it does not comply with the conditions of use. If conditions of use are changed, your PHR provider is required to notify you of the changes.

CCHIT4. Amending the Record. You should have the ability to change or request changes to your health record via email or telephone, and the telephone number of customer service must be posted on the Web site of your PHR provider.

CCHIT5. Account Management. Your PHR provider must have a way for you to terminate your account, if you wish, and to confirm that all your personal data has been deleted from the system.

CCHIT6. Document Import. Your PHR system should be able to retrieve health records, explicitly label and manage your personal health information and be able to distinguish between data entered by you and data retrieved from other sources.

CCHIT7. Data Availability. Your system should allow you to view or print your health information whenever you need it.

It is instructive to compare CCHIT criteria with the HPP Best Practices for PHR systems, since both attempt to cover the same ground. CCHIT1 (consent) is covered by BP3 (choose which content), BP4 (control access), and BP1 (transparency). CCHIT2 (control) is similar to BP4 (control). CCHIT3 (conditions of use) is similar to BP1 (transparency and notice). CCHIT4 (amending) is mentioned by BP3, although BP3 calls it "annotating" rather than "amending", which in some contexts may be an important difference. CCHIT5 (the ability to delete your account) is not covered by any HPP Best Practice, which is an interesting omission. CCHIT6 (document import) is about capability (ability to import documents) and about distinguishing user-entered data from provider data (which is mentioned by BP3). CCHIT7 (availability) is not covered by any HPP Best Practice, even BP8 (data management).

There are many of HPP's Best Practices that do not appear in the CCHIT criteria. In particular, BP1 is more specific than CCHIT3 about the presentation and quality of the terms and notification, BP2 (education of the user) is not covered, BP3 is more clear that sources of data should be clearly identified, BP4 precludes employer use of employee PHR data, BP4 requires an audit trail and that the audit trail be easily available, BP5 allows the designation of proxies, BP6 requires the "chain of trust" in which privacy policies follow the data to other business partners, BP7 requires strong mechanisms for security, access control, and access logs, BP8 requires data-integrity policies and mechanisms, BP9 requires notice and assistance in the event of a data breach, and BP10 requires PHR portability to new employers. None of these aspects— many of which are important properties—are covered by the CCHIT criteria, leading us to the opinion that the CCHIT Criteria are too weak and should be strengthened.

Similarly, the CCHIT criteria fall short of the ONC framework. ON3 (openness and transparency) is covered somewhat by CCHIT3 (conditions of use), but not the requirement (stated in the details) for access to an audit log. ONC4 (individual choice) says (in the details) that a Patient should be able to designate a proxy, which is not mentioned in the Criteria. ONC5 (collection, use, and disclosure) limits the collection, use, and disclosure of PHI to the minimum necessary for a particular purpose and the Criteria have no such concept; of course, a PHR is intended to collect PHI more widely than an EHR, at the discretion of the Patient, but this concept should still apply to the disclosure and use of PHR. ONC6 (data quality and integrity) states the obligation of providers to ensure records are complete and correct; no such requirement appears in the Criteria. ONC7 (safeguards) describes security and data-integrity requirements, which do not appear in the Criteria. Although CCHIT1 says the Patient should be able to challenge the PHR provider if they do not follow their terms of use, the Criteria have nothing like ONC8 (accountability) that requires the PHR provider to have mechanisms for monitoring internal compliance, and for notifying Patients if there is a data breach.

A.4.1. Others. There are other privacy frameworks worth a brief mention.

The Organization for Economic Co-operation and Development (OECD) Guidelines on the Protection of Privacy and Transborder Flows of Personal Data [OECD 1980] is an old (1980) list of eight principles that, although not specific to healthcare, has inspired all of these frameworks.

Buckovich et al. [1999] prepared a detailed survey of privacy principles for healthcare in 1998. Although interesting, this survey pre-dates all of the healthcare documents we survey here, even pre-dating HIPAA

The Health Information Trust Alliance (HITRUST) released a security framework for healthcare in March 2009, but it is available only to member organizations and only for a fee [Kaplan 2009]. Little information is available about this framework and it is not clear whether it includes a comprehensive set of privacy principles.

Two international documents address privacy principles, though not specifically for healthcare. In 2005, the Asia-Pacific Economic Community (APEC) published a privacy framework, a set of "information privacy principles" [APEC 2005]. This framework covers the commercial use of personal information. One analysis shows, however, that there are too few details—and far too much flexibility—in the APEC Framework for it to be useful [Pounder 2007]. In 2007, the International Security, Trust, and Privacy Alliance compared several important privacy frameworks from around the world [ISTPA 2007]. This document brings a broad, international perspective to privacy frameworks, but its focus is on providing uniform definitions for common terms; it does little, specifically, to generate a precise set of principles.

A.5. Summary

There are several existing conceptual privacy frameworks. Each has been developed by experts in the field, usually by large bodies of diverse stakeholders who have worked for months or years to examine relevant laws, documents, and current practices. Many of the frameworks have been developed after studying the others, and thus there is a significant degree of overlap. On the other hand, there is a surprising degree of difference in the principles, and some frameworks have notable omissions. There is fairly broad agreement about general principles, but there is room for reasonable disagreement about the details, in part because there are difficult trade-offs between protecting privacy and providing efficient, effective healthcare. (For example, consider a patient that has configured access to his medical record so that access to data about his sexual diseases and psychological troubles are hidden from all but a few doctors; if he presents with an emergency condition then the lack of that information may lead to inefficient or incorrect care. Or, consider a doctor who gives her password to the nurse to avoid the inconvenience of logging herself in every time she needs to update a prescription; this may lead to more efficient care but increases risk to patient privacy. Or, consider mobile sensor nodes that are easy to pair with the Patient's mobile node, but this ease of use is based on a weak pairing protocol that is vulnerable to man-in-the-middle attacks.)

The NCVHS discussed several of these issues in 2006 [Cohn 2006], and revisited some of them in 2008 [NCVHS 2008].

APPENDIX
B. GLOSSARY OF ACRONYMS AND TERMS

AHIC	American Health Information Community, a federal advisory body chartered to make recommendations to HHS on accelerating the development and adoption of health information technology.
APEC	Asia-Pacific Economic Community.
ARRA	American Recovery and Reinvestment Act of 2009. Extends HIPAA's privacy regulations.
CCHIT	Certification Commission for Healthcare Information Technology.
CDT	Center for Democracy and Technology.

Client	The Consumer's computer, in our conceptual architecture.
Common Framework	The Markle Foundation launched a project "Connecting for Health", which brought together a wide range of stakeholders in developing a "Common Framework," a model for healthcare information exchange [MF 2008].
Conceptual privacy framework	A coherent set of actionable principles to protect Patients' health information privacy.
Confidentiality	"The obligations of those who receive information to respect the privacy interests of those to whom the data relate" [Cohn 2006]. See also *privacy* and *security*.
Consumer	The user of mHealth sensor data, in our conceptual architecture. We capitalize it as a reminder that "Consumer" includes a broad range of users of mHealth data, from clinical staff to wellness coaches to relatives.
Covered Entities	Organizations that fall under HIPAA regulations.
Ecosystem	The set of organizations and agencies that support the standardization, manufacture, certification, distribution, management, and operation of mHealth systems and services.
EHR	Electronic Health Record. "An electronic record of health related information on an individual that conforms to nationally recognized interoperability standards and that can be created, managed, and consulted by authorized clinicians and staff across more than one health care organization" [NAHIT 2008]. Contrast with *EMR* and *PHR*.
EMR	Electronic Medical Record. "An electronic record of health related information on an individual that can be created, gathered, managed, and consulted by authorized clinicians and staff within one health care organization" [NAHIT 2008]. Contrast with *EHR*.
Health information privacy	"An individual's right to control the acquisition, uses, or disclosures of his or her identifiable health data" [Cohn 2006]. See also *confidentiality*.
HHS	US Department of Health and Human Services.
HIE	Health Information Exchange.
HIPAA	Health Information Portability and Accountability Act of 1996. US law specifying, in part, privacy rules regarding Protected Health Information and the responsibilities of Covered Entities [Scholl et al. 2008]. See also *HITECH*.
HITECH	Health Information Technology for Economic and Clinical Health Act of 2009 [HITECH1 2009; HITECH2 2009]. See also *HIPAA*.

HPP	Health Privacy Project [HPP 2007], recently adopted by the Center for Democracy & Technology (CDT).
HRS	Health Records System, in our conceptual architecture.
ISTPA	International Security, Trust, and Privacy Alliance.
Markle Foundation	an organization concerned with technology, healthcare, and national security. See also *Common Framework*.
mHealth	The use of mobile computing and communications technology in the delivery of healthcare or collection of health information. The term "mHealth" applies broadly to the use of mobile technology in healthcare applications. In this document, however, we focus on patient-centered technology.
MID	Mobile Internet Device, a multimedia-capable handheld computer providing wireless Internet access [MID 2009]. MIDs are larger than smart phones but smaller than netbooks and laptops.
MN	Mobile Node, in our conceptual architecture.
NCVHS	National Committee on Vital Health and Statistics (US).
NHIN	Nationwide Health Information Network, a US effort to provide a secure, nationwide, interoperable health information infrastructure that will connect providers, consumers, and others involved in supporting health and healthcare.
NHS	National Health Service in the UK.
ONC	Office of the National Coordinator for Health Information Technology (in the US Department of Health and Human Services).
Patient	The subject of mHealth sensing, in our conceptual architecture. We capitalize it as a reminder that a "Patient" is not restricted to a "patient" in the clinical sense, although we use "patient" when we are indeed in a clinical context.
PHI	Personal Health Information. Any personally identifiable information (PII) that is health-related. We follow NCVHS and define PHI as "personal health information" rather than "protected health information", which is a phrase that has specific meaning in a specific legal context (HIPAA) [Cohn 2006]. HIPAA defines "protected health information" as individually identifiable health information that is stored or transmitted in any medium, including information related

| | to the Patient's health, the provision or healthcare, or billing for healthcare. Contrast with *PII* and *PrHI*. |

PHR — Personal Health Record. "An electronic record of health related information on an individual that conforms to nationally recognized interoperability standards and that can be drawn from multiple sources while being managed, shared, and controlled by the individual" [NAHIT 2008]. Contrast with *EHR*.

PII — Personally Identifiable Information. Any information that uniquely identifies an individual, such as a name, social security number, patient ID number, or street address. Sometimes a MAC address or IP address is considered personally identifiable information. Contrast with *PHI*.

PKI — Public Key Infrastructure. A set of services and protocols that support the use of public-key cryptography, including certificate authorities (who generate and verify certificates that bind keys to identity or attributes).

PRBAC — Privacy-aware RBAC, see RBAC.

PrHI — Protected Health Information. A term specific to HIPAA, Protected Health Information is PHI that is covered by HIPAA privacy rules. PrHI is our abbreviation, to distinguish it from *PHI*.

Privacy — see *Health information privacy*.

RBAC — Role-based Access Control.

RHIO — Regional Health Information Organization.

Security — "Physical, technological, or administrative safeguards or tools used to protect identifiable health data from unwarranted access or disclosure" [Cohn 2006]. See also *privacy* and *confidentiality*.

SN — Sensor Node, in our conceptual architecture.

TPM — Trusted Platform Module, as defined by the Trusted Computing Group [TPM 2009].

ACKNOWLEDGMENTS

REFERENCES

ACKERMAN, M. S. AND MAINWARING, S. D. 2005. Privacy issues and human-computer interaction. In *Security and Usability: Designing Secure Systems that People Can Use*, L. F. Cranor and S. Garfinkel, Eds., O'Reilly Media, 381–400. http://oreilly.com/catalog/9780596008277/.

ACLU 2009, American Civil Liberties Union. The American Recovery and Reinvestment Act of 2009: Health information technology, privacy summary. http://www.aclu.org/images/asset_upload_file625_38771.pdf. (last accessed 3/09).

AGRAFIOTI, F. AND HATZINAKOS, D. 2008. Fusion of ECG sources for human identification. In *Proceedings of the International Symposium on Communications, Control and Signal Processing (ISCCSP)*. IEEE Press, 1542–1547. DOI 10.1109/ISCCSP.2008.4537472.

AL AMEEN, M., LIU, J., AND KWAK, K. 2010. Security and privacy issues in wireless sensor networks for healthcare applications. *J. Medical Syst.* 1–9. DOI 10.1007/s10916-010-9449-4.

ALLONE HEALTH. 2009. PHR access on mobile phone. http://www.allonemobile.com. (last accessed 3/09)

AMERICAN MEDICAL ASSOCIATION. 2009. HR.1, the American Recovery and Reinvestment Act of 2009: Explanation of privacy provisions. http://www.ama-assn.org/ama1/pub/upload/mm/399/arra-privacy-provisions.pdf. (last accessed 3/09).

ANDERSEN, J. 2009. Secure group formation protocol for a medical sensor network prototype. In *Proceedings of the International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*. IEEE, 343–348. DOI 10.1109/ISSNIP.2009.5416771.

ANVITA HEALTH. 2009. Google health on mobile phone. http://www.anvitahealth.com. (last accessed 3/09).

APEC 2005. APEC privacy framework. http://tinyurl.com/cusnax.

APPARI, A. AND JOHNSON, M. E. 2010. Information security and privacy in healthcare: Current state of research. *Int. J. Internet Enterprise Manage.* 6, 4, 279–314. http://mba.tuck.dartmouth.edu/pages/faculty/eric.johnson/pdfs/AJIJIEM.pdf.

AYLWARD, R. AND PARADISO, J. A. 2007. A compact, high-speed, wearable sensor network for biomotion capture and interactive media. In *Proceedings of the International Workshop on Information Processing in Sensor Networks (IPSN)*. ACM, 380–389. DOI 10.1145/1236360.1236408.

BAKER, C. R., ARMIJO, K., BELKA, S., BENHABIB, M., BHARGAVA, V., BURKHART, N., DER MINASSIANS, A., DERVISOGLU, G., GUTNIK, L., HAICK, B. M., HO, C., KOPLOW, M., MANGOLD, J., ROBINSON, S., ROSA, M., SCHWARTZ, M., SIMS, C., STOFFREGEN, H., WATERBURY, A., LELAND, E. S., PERING, T., AND WRIGHT, P. K. 2007. Wireless sensor networks for home health care. In *Proceedings of the International Conference on Advanced Information Networking and Applications Workshops*. IEEE Computer Society, 832–837. DOI 10.1109/AINAW.2007.376.

BALDUS, H., KLABUNDE, K., AND MÜSCH, G. 2004. Reliable set-up of medical body-sensor networks. In *Proceedings of the 1st European Workshop on Wireless Sensor Networks*. Lecture Notes in Computer Science, vol. 2920. Springer, 353–363. DOI 10.1007/978-3-540-24606-0-24.

BARTH, A., DATTA, A., MITCHELL, J. C., AND NISSENBAUM, H. 2006. Privacy and contextual integrity: Framework and applications. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*. IEEE Press, 15–29. DOI 10.1109/SP.2006.32.

BARTH, A. T., HANSON, M. A., POWELL, H. C., UNLUER, D., WILSON, S. G., AND LACH, J. 2008. Body-coupled communication for body sensor networks. In *Proceedings of the ICST International Conference on Body Area Networks (BodyNets)*. Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering (ICST), 1–4. Online at http://portal.acm.org/citation.cfm?id=1460257.1460273.

BECHER, E., BENENSON, Z., AND DORNSEIF, M. 2006. Tampering with motes: Real-world physical attacks on wireless sensor networks. In *Proceedings of the International Conference on Security in Pervasive Computing (SPC)*. Springer-Verlag, 104–118. DOI 10.1007/11734666 9.

BEKIARIS, E., DAMOUSIS, I. G., AND TZOVARAS, D. 2008. Unobtrusive multimodal biometric authentication: The HUMABIO project concept. *EURASIP J. Adv. Sig. Process.* DOI 10.1155/2008/265767.

BELLMAN, S., JOHNSON, E. J., AND LOHSE, G. L. 2001. To opt-in or opt-out? it depends on the question. *Comm. ACM 44*, 2, 25–27. DOI 10.1145/359205.359241.

BICHLER, D., STROMBERG, G., HUEMER, M., AND LÖW, M. 2007. Key generation based on acceleration data of shaking processes. In *Proceedings of Ubiquitous Computing (UbiComp)*. Lecture Notes in Computer Science Series, vol. 4717. Springer-Verlag, 304–317. DOI 10.1007/978-3-540-74853-3 18.

BLOUGH, D., AHAMAD, M., LIU, L., AND CHOPRA, P. 2008. MedVault: Ensuring security and privacy for electronic medical records. NSF CyberTrust Principal Investigators Meeting. Online at http://www.cs.yale.edu/cybertrust08/posters/posters/158 medvault_poster_CT08.pdf.

BORIC-LUBECKE, O. AND LUBECKE, V. M. 2002. Wireless house calls: using communications technology for health care and monitoring. *IEEE Microwave Magazine 3*, 3, 43–48. DOI 10.1109/MMW.2002.1028361.

BRAHMBHATT, B. 2010. Position and perspective of privacy laws in India. In *AAAI Spring Symposium Series: Intelligent Information Privacy Management*. AAAI. Online at http://www.aaai.org/ocs/index.php/SSS/SSS10/paper/view/1197/1474.

BRATUS, S., CORNELIUS, C., KOTZ, D., AND PEEBLES, D. 2008. Active behavioral fingerprinting of wireless devices. In *Proceedings of the ACM Conference on Wireless Network Security (WiSec)*. ACM, 56–61. DOI 10.1145/1352533.1352543.

BREAUX, T. D. AND ANTÓN, A. I. 2008. Analyzing regulatory rules for privacy and security requirements. *IEEE Trans. Softw. Eng.* 34, 1, 5–20. DOI 10.1109/TSE.2007.70746.

BRIK, V., BANERJEE, S., GRUTESER, M., AND OH, S. 2008. Wireless device identification with radiometric signatures. In *Proceedings of the ACM International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 116–127. DOI 10.1145/1409944.1409959.

BUCKOVICH, S. A., RIPPEN, H. E., AND ROZEN, M. J. 1999. Driving toward guiding principles: A goal for privacy, confidentiality, and security of health information. *J. AMIA 6*, 2, 122–133. DOI 10.1136/jamia.1999.0060122.

CCHIT 2008. Consumer's guide to certification of personal health records. Booklet. Online at http://cchit.org/files/CCHITPHRConsumerGuide08.pdf.

CDT 2008. Comprehensive privacy and security: Critical for health information technology. White paper. Online at http://www.cdt.org/healthprivacy/20080514HPframe.pdf.

CDT 2009. Summary of health privacy provisions in the 2009 economic stimulus legislation. White paper. Online at http://www.cdt.org/healthprivacy/20090324_ARRAPrivacy.pdf.

CHCF 2008. Whose data is it anyway? Expanding consumer control over personal health information. California Healthcare Foundation. Online at http://ehealth.chcf.org/topics/view.cfm?itemID=133577.

CHERUKURI, S., VENKATASUBRAMANIAN, K. K., AND GUPTA, S. K. S. 2003. BioSec: A biometric based approach for securing communication in wireless networks of biosensors implanted in the human body. In *Proceedings of the International Conference on Parallel Processing Workshops*. IEEE Computer Society, 432–439. DOI 10.1109/ICPPW.2003.1240399.

CHOI, Y. B., CAPITAN, K. E., KRAUSE, J. S., AND STREEPER, M. M. 2006. Challenges associated with privacy in healthcare industry: Implementation of HIPAA and security rules. *J. Med. Syst. 30*, 1, 57–64. DOI 10.1007/s10916-006-7405-0.

COHN, S. P. 2006. Privacy and confidentiality in the nationwide health information network. Online at http://www.ncvhs.hhs.gov/060622lt.htm.

COLLINS, T. 2006. NHS trust uncovers password sharing risk to patient data. *Computer Weekly*. Online at http://www.computerweekly.com/Articles/2006/07/11/216882/nhs-trust-uncovers-password-sharing-risk-to-patient.htm.

CORNELIUS, C., AND KOTZ, D. 2010. On usable authentication for wireless body area networks. In *Proceedings of the USENIX Workshop on Health Security and Privacy*. USENIX Association. Online at http://www.cs.dartmouth.edu/_dfk/papers/abstracts/cornelius-healthsec10.html.

CORNELIUS, C. AND KOTZ, D. 2011. Recognizing whether sensors are on the same body. In *Proceedings of the International Conference on Pervasive Computing. Lecture Notes in Computer Science*. Springer, 332–349. DOI 10.1007/978-3-642-21726-5 21.

COVINGTON, M., MOYER, M., AND AHAMAD, M. 2000. Generalized role-based access control for securing future applications. In *Proceedings of the National Information Systems Security Conference*. NIST. Online at http://csrc.nist.gov/nissc/2000/proceedings/papers/040.pdf.

CRANOR, L. F. 2003. 'I didn't buy it for myself': Privacy and ecommerce personalization. In *Proceedings of the ACM Workshop on Privacy in the Electronic Society (WPES)*. ACM, 111–117. DOI 10.1145/1005140.1005158.

CRANOR, L. F. 2005. Privacy policies and privacy preferences. In *Security and Usability: Designing Secure Systems that People Can Use*. L. F. Cranor and S. Garfinkel, Eds. O'Reilly Media, Chapter 22, 447–469. Online at http://oreilly.com/catalog/9780596008277/.

CRANOR, L. F. 2008. A framework for reasoning about the human in the loop. In *Proceedings of the Conference on Usability, Psychology, and Security (UPSEC)*. USENIX Association, 1–15. Online at http://static.usenix.org/event/upsec08/tech/full_pasess/cranor/cranor.pdf.

DAI ZOVI, D. A. AND MACAULAY, S. A. 2005. Attacking automatic wireless network selection. In *Proceedings of the IEEE SMC Information Assurance Workshop*. IEEE Press, 365–372. DOI 10.1109/IAW.2005.1495975.

DE MULDER, Y., DANEZIS, G., BATINA, L., AND PRENEEL, B. 2008. Identification via location-profiling in GSM networks. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES)*. ACM, 23–32. DOI 10.1145/1456403.1456409.

DH 2008, Intel Research. Digital Home project. Online at http://www.intel.com/research/exploratory/digitalhome.htm, visited Mar. 2008.

DIT 2011, Government of India, Department of Information Technology (DIT). Information Technology Act 2000 [India]. Online at http://www.mit.gov.in/content/information-technology-act-2000, visited Feb. 2011.

DIXON, P. 2006. Medical identity theft: The information crime that can kill you. Online at http://www.worldprivacyforum.org/pdf/wpfmedicalidtheft2006.pdf.

DOMINGO-FERRER, J., MARTÍNEZ-BALLESTÉ, A., MATEO-SANZ, J. M. AND SEBÉ, F. 2006. Efficient multivariate data-oriented microaggregation. *VLDB J. 15*, 4, 355–369. DOI 10.1007/s00778-006-0007-0.

DS 2009, Daily Strength. Dailystrength.org. Online at http://www.dailystrength.org/, visited Oct. 2009.

EISENMAN, S. B., MILUZZO, E., LANE, N. D., PETERSON, R. A., AHN, G.-S., AND CAMPBELL, A. T. 2009. BikeNet: A mobile sensing system for cyclist experience mapping. *ACM Trans. Sensor Netw. (TOSN) 6*, 1, 1–39. DOI http://doi.acm.org/10.1145/1653760.1653766.

ENCK, W., ONGTANG, M., AND MCDANIEL, P. 2009. On lightweight mobile phone application certification. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM, 235–245. DOI 10.1145/1653662.1653691.

EU 2009, Office of the Data Protection Commissioner. EU Directive 95/46/EC: The data protection directive. Online at http://www.dataprotection.ie/viewdoc.asp?DocID=92, visited Mar. 2009.

FERRAIOLO, D. AND KUHN, R. 1992. Role based access control. In *Proceedings of the National Computer Security Conference*. NIST. Online at http://csrc.nist.gov/rbac/ferraiolo-kuhn-92.pdf.

FRANK, M., STREICH, A. P., BASIN, D., AND BUHMANN, J. M. 2009. A probabilistic approach to hybrid role mining. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM, 101–111. DOI 10.1145/1653662.1653675.

FRANKLIN, J., MCCOY, D., TABRIZ, P., NEAGOE, V., RANDWYK, J. V., AND SICKER, D. 2006. Passive data link layer 802.11 wireless device driver fingerprinting. In *Proceedings of the USENIX Security Symposium*. USENIX Association, 167–178. Online at http://www.usenix.org/events/sec06/tech/franklin.html.

FRIEDMAN, B., LIN, P., AND MILLER, J. K. 2005. *Informed consent by design. In Security and Usability: Designing Secure Systems that People Can Use*. L. F. Cranor and S. Garfinkel, Eds. O'Reilly Media, Chapter 24, 495–521. Online at http://oreilly.com/catalog/9780596008277/.

GARCIA-MORCHON, O. AND BALDUS, H. 2008. Efficient distributed security for wireless medical sensor networks. In *Proceedings of the International Conference on Intelligent Sensors, Sensor Networks and Information Processing*. IEEE, 249–254. DOI 10.1109/ISSNIP.2008.4761995.

GARCIA-MORCHON, O., FALCK, T., HEER, T., AND WEHRLE, K. 2009. Security for pervasive medical sensor networks. In *Proceedings of the International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQuitous)*. IEEE Press. DOI 10.4108/ICST.MOBIQUITOUS2009.6832.

GD. 2011. Giesecke and Devrient GmbH. Online at http://www.gi-de.com/, visited Mar. 2011.

GEDIK, B. AND LIU, L. 2008. Protecting location privacy with personalized k-anonymity: Architecture and algorithms. *IEEE Trans. Mobile Comput. 7*, 1, 1–18. DOI 10.1109/TMC.2007.1062.

GH 2008, Google. Google Health. Online at https://www.google.com/health, visited Nov. 2008.

GEORGIA INSTITUTE OF TECHNOLOGY. 2008. Aware Home project. http://www.cc.gatech.edu/fce/ahri/. (last accessed 3/08).

GIANNETSOS, T., DIMITRIOU, T., AND PRASAD, N. R. 2011. People-centric sensing in assistive healthcare: Privacy challenges and directions. *Secur. Commun. Netw*. DOI 10.1002/sec.313.

GILBERT, P., COX, L. P., JUNG, J., AND WETHERALL, D. 2010. Toward trustworthy mobile sensing. In *Proceedings of the Workshop on Mobile Computing Systems & Applications (HotMobile)*. ACM, 31–36. DOI 10.1145/1734583.1734592.

GOLDMAN, J. 1998. Protecting privacy to improve health care. *Health Affairs 17*, 6, 47–60. DOI 10.1377/hlthaff.17.6.47.

GOLLE, P. AND PARTRIDGE, K. 2009. On the anonymity of home/work location pairs. In *Proceedings of Pervasive Computing*. Lecture Notes in Computer Science Series, vol. 5538. Springer-Verlag, 390–397. DOI 10.1007/978-3-642-01516-8_26.

GOYAL, V. 2007. Certificate revocation using fine grained certificate space partitioning. In *Proceedings of the International Conference on Financial Cryptography and Data Security (FCDS)*. S. Dietrich and R. Dhamija, Eds. Lecture Notes in Computer Science Series, vol. 4888. Springer-Verlag, 247–259. DOI 10.1007/978-3-540-77366-5_24.

GOYAL, V., PANDEY, O., SAHAI, A., AND WATERS, B. 2006. Attribute-based encryption for fine-grained access control of encrypted data. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM, 89–98. DOI 10.1145/1180405.1180418.

GREENSTEIN, B., MCCOY, D., PANG, J., KOHNO, T., SESHAN, S., AND WETHERALL, D. 2008. Improving wireless privacy with an identifier-free link layer protocol. In *Proceedings of the International Conference on Mobile Systems, Applications and Services (MobiSys)*. ACM Press, 40–53. DOI 10.1145/1378600.1378607.

GRUTESER, M. AND GRUNWALD, D. 2003. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proceedings of the International Conference on Mobile Systems, Applications and Services (MobiSys)*. ACM, 31–42. DOI 10.1145/1066116.1189037.

GUTMANN, P. 2002. PKI: It's not dead, just resting. *IEEE Computer 35*, 8, 41–49. DOI 10.1109/ MC.2002.1023787.

HALAMKA, J. 2008. Respecting patient privacy preferences. Blog– Life as a Healthcare CIO. Online at http://geekdoctor.blogspot.com/2008/ 01/respecting-patient-privacy-preferences.html.

HALAMKA, J., LEAVITT, M., AND TOOKER, J. 2009. A shared roadmap and vision for health IT. Position statement. Online at http://tinyurl.com/c8ztuy.

HALPERIN, D., HEYDT-BENJAMIN, T. S., RANSFORD, B., CLARK, S. S., DEFEND, B., MORGAN, W., FU, K., KOHNO, T., AND MAISEL, W. H. 2008a. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*. IEEE Press, 129–142. DOI 10.1109/SP.2008.31.

HALPERIN, D., THOMAS, FU, K., KOHNO, T., AND MAISEL, W. H. 2008b. Security and privacy for implantable medical devices. *IEEE Pervas. Comput.* 7, 1, 30–39. DOI 10.1109/MPRV.2008.16.

HHS 2009, US Department of Human and Health Services. Draft model personal health record (PHR) privacy notice & facts-at-a-glance. Online at http://tinyurl.com/cxm4q3, visited Apr. 2009.

HIPAA 2010, HHS. HIPAA website. Online at http://www.hhs.gov/ocr/privacy/, visited Mar. 2010.

HITECH1 2009, Coppersmith Gordon Schermer and Brockelman. HITECH Act expands HIPAA privacy and security rules. Online at http://www.azhha.org/member_and_media_resources/documents/HITECHAct. pdf, visited Nov. 2009.

HITECH2 2009, HIPAA Survival Guide. HITECH Act text. Online at http://www.hipaasurvivalguide.com/ hitech-act-text.php, visited Nov. 2009.

HITSP 2008. TP-30: HITSP manage consent directives transaction package. Online at http://www.hitsp.org/ ConstructSet Details.aspx?&PrefixAlpha=2&PrefixNumeric=30.

HL 2009, Health Law News and Notes. FAQs on ARRA/Stimulus Bill changes for business associates. Online at http://healthlawoffices. com/blog/?p=85, visited Mar. 2009.

HOH, B. AND GRUTESER, M. 2005. Protecting location privacy through path confusion. In *Proceedings of the IEEE/CreateNet International Conference on Security and Privacy for Emerging Areas in Communication Networks (SecureComm)*. IEEE Press. DOI 10.1109/SECURECOMM.2005.33.

HOH, B., GRUTESER, M., XIONG, H., AND ALRABADY, A. 2007. Preserving privacy in GPS traces via uncertainty-aware path cloaking. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM, 161–171. DOI 10.1145/1315245.1315266.

HPP 1999. Best principles for health privacy. Georgetown University. Online at http://www.healthprivacy. org/usr doc/33807.pdf.

HPP 2007. Best practices for employers offering personal health records (PHRs). Developed by the Employers' Working Group on Personal Health Records (PHRs). Online at http://www.cdt.org/healthprivacy/ 2007Best Practices.pdf.

HU, W., CORKE, P., SHIH, W. C., AND OVERS, L. 2009. secFleck: A public key technology platform for wireless sensor networks. In *Proceedings of the European Conference on Wireless Sensor Networks (EWSN)*. Springer-Verlag, 296–311. DOI 10.1007/978-3-642-00224-3 19.

IACHELLO, G. AND HONG, J. 2007. End-user privacy in human-computer interaction. *Found. Trends Hum.-Comput. Interact. (FTHCI)* 1, 1–137. DOI 10.1561/1100000004.

IHE 2009, IHE International. IHE profiles. Online at http://www.ihe.net/profiles/index.cfm, visited Nov. 2009.

INDIA 2011. Information technology rules GSR 313(E)-316(E). Government of India. http://deity.gov.in/sites/ upload_files/dit/files/GSR3_10511(1).pdf.

IRVINE, J. M., ISRAEL, S. A., SCRUGGS, T. W., AND WOREK, W. J. 2008. eigenPulse: Robust human identification from cardiovascular function. *Patt. Recog. 41*, 11, 3427–3435. DOI 10.1016/j.patcog. 2008.04.015.

ISTPA 2007. Analysis of privacy principles: Making privacy operational. Online at http://www.istpa. org/pdfs/ISTPAAnalysisofPrivacyPrinciplesV2.pdf.

JAIN, A. K., FLYNN, P., AND ROSS, A. A., EDS. 2007. Handbook of Biometrics. Springer-Verlag. Online at http://www.springer.com/computer/computer+imaging/book/978-0-387-71040-2.

JAIN, A. K., ROSS, A., AND PRABHAKAR, S. 2004. An introduction to biometric recognition. *IEEE Trans. Circ. Syst. Video Tech. 14*, 1, 4–20. DOI 10.1109/TCSVT.2003.818349.

JANA, S., PREMNATH, S. N., CLARK, M., KASERA, S. K., PATWARI, N., AND KRISHNAMURTHY, S. V. 2009. On the effectiveness of secret key extraction from wireless signal strength in real environments. In *Proceedings of the ACM International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 321–332. DOI 10.1145/1614320.1614356.

JEA, D., LIU, J., SCHMID, T., AND SRIVASTAVA, M. B. 2008. Hassle free fitness monitoring. In *Proceedings of the Workshop on Systems and Networking Support for Healthcare and Assisted Living Environments (HealthNet)*. ACM. DOI 10.1145/1515747.1515756.

JOHNSON, M. E. 2009. Data hemorrhages in the health-care sector. In *Financial Cryptography and Data Security*. Springer-Verlag. DOI 10.1007/978-3-642-03549-4_5.

JONES, V., MEI, H., BROENS, T., WIDYA, I., AND PEUSCHER, J. 2007. Context aware body area networks for telemedicine. In *Advances in Multimedia Information Processing (PCM)*. Springer-Verlag, 590–599. DOI 10.1007/978-3-540-77255-2_74.

KAPLAN, D. 2009. Group unveils first-of-its-kind standard to secure patient data. SC Magazine. Online at http://www.scmagazineus.com/Group-unveils-first-of-its-kind-standard-to-secure-patient-data/article/128168/.

KARAT, C., BRODIE, C., AND KARAT, J. 2005. Usability design and evaluation for privacy and security solutions. In *Security and Usability: Designing Secure Systems that People Can Use*, L. F. Cranor and S. Garfinkel, Eds. O'Reilly Media, Chapter 4, 47–74. Online at http://oreilly.com/catalog/9780596008277/.

KELLEY, P. G., CESCA, L., BRESEE, J., AND CRANOR, L. F. 2010. Standardizing privacy notices: an online study of the nutrition label approach. In *Proceedings of the International Conference on Human Factors in Computing Systems (CHI)*. ACM, 1573–1582. DOI 10.1145/1753326.1753561.

KELLEY, P. G., HANKES DRIELSMA, P., SADEH, N., AND CRANOR, L. F. 2008. User-controllable learning of security and privacy policies. In *Proceedings of the ACM Workshop on Security and Artificial Intelligence (AIsec)*. ACM, 11–18. DOI 10.1145/1456377.1456380.

KLASNJA, P., CONSOLVO, S., CHOUDHURY, T., AND BECKWITH, R. 2009. Exploring privacy concerns about personal sensing. In *Proceedings of the International Conference on Pervasive Computing (Pervasive)*. Springer-Verlag. DOI 10.1007/978-3-642-01516-8_13.

KLEIDERMACHER, D. 2008. Next generation secure mobile devices. *Inf. Quart. 7*, 4, 14–17. Online at http://www.iqmagazineonline.com/article.php?issue=25&article_id=1041.

KOTZ, D. 2011. A threat taxonomy for mHealth privacy. In *Proceedings of the Workshop on Networked Healthcare Technology (NetHealth)*. IEEE Press. DOI 10.1109/COMSNETS.2011.5716518.

KOTZ, D., AVANCHA, S., AND BAXI, A. 2009. A privacy framework for mobile health and home-care systems. In *Proceedings of the Workshop on Security and Privacy in Medical and Home-Care Systems (SPIMACS)*. ACM, 1–12. DOI 10.1145/1655084.1655086.

KUIE, T. S. 2003. The impact of data privacy protection in medical practice in Singapore. *SGH Proc. 12*, 4, 201–207. Online at http://www.pgmi.com.sg/SGHproceeding/12-4/impact%20of%20data%20privacy.pdf.

KULKARNI, P. AND ÖZTÜRK, Y. 2007. Requirements and design spaces of mobile medical care. *SIGMOBILE Mobile Comput. Commun. Rev. 11*, 3, 12–30. DOI 10.1145/1317425.1317427.

KUMAR, A., SAXENA, N., TSUDIK, G., AND UZUN, E. 2009. A comparative study of secure device pairing methods. *Pervas. Mobile Comput. 5*, 6, 734–749. DOI 10.1016/j.pmcj.2009.07.008.

KUMARAGURU, P. AND CRANOR, L. 2006. Privacy in India: Attitudes and awareness. In *Proceedings of the International Workshop on Privacy Enhancing Technologies (PET)*, G. Danezis and D. Martin, Eds. Springer, 243–258. DOI 10.1007/11767831_16.

LIU, A. AND NING, P. 2008. TinyECC: A configurable library for elliptic curve cryptography in wireless sensor networks. In *Proceedings of the International Workshop on Information Processing in Sensor Networks (IPSN)*. IEEE Press. DOI 10.1109/IPSN.2008.47.

LOWRANCE, W. W. 2009. Privacy and health research: New laws in Europe. The HHS Data Council, US Department of Health and Human Services. Online at http://aspe.hhs.gov/datacncl/PHR5.htm.

MACHANAVAJJHALA, A., GEHRKE, J., KIFER, D., AND VENKITASUBRAMANIAM, M. 2006. *l*-diversity: Privacy beyond *k*-anonymity. In *Proceedings of the International Conference on Data Engineering (ICDE)*. IEEE Press, 24–85. DOI 10.1109/ICDE.2006.1.

MACK, D. C., ALWAN, M., TURNER, B., SURATT, P., AND FELDER, R. A. 2006. A passive and portable system for monitoring heart rate and detecting sleep apnea and arousals: Preliminary validation. In *Proceedings of the Transdisciplinary Conference on Distributed Diagnosis and Home Healthcare (D2H2)*. IEEE Computer Society, 51–54. DOI 10.1109/DDHH.2006.1624795.

MALAN, D. J., WELSH, M., AND SMITH, M. D. 2008. Implementing public-key infrastructure for sensor networks. *ACM Trans. Sensor Netw. (TOSN) 4*, 4, 1–23. DOI 10.1145/1387663.1387668.

MALASRI, K. AND WANG, L. 2007. Addressing security in medical sensor networks. In *Proceedings of the Workshop on Systems and Networking Support for Healthcare and Assisted Living Environments (HealthNet)*. ACM Press, 7–12. DOI 10.1145/1248054.1248058.

MALASRI, K. AND WANG, L. 2008. Design and implementation of a secure wireless mote-based medical sensor network. In *Proceedings of Conference on Ubiquitous Computing (UbiComp)*. ACM, 172–181. DOI 10.1145/1409635.1409660.

MALIN, B. 2006. Re-identification of familial database records. In *Proceedings of the AMIA Annual Symposium*. AMIA, 524–528. Online at http://view.ncbi.nlm.nih.gov/pubmed/17238396.

MALIN, B. AND AIROLDI, E. 2007. Confidentiality preserving audits of electronic medical record access. *Stud. Health Tech. Informat.* 129, Part 1, 320–324. Online at http://view.ncbi.nlm.nih. gov/pubmed/17911731.

MARE, S. AND KOTZ, D. 2010. Is Bluetooth the right technology for mHealth? In USENIX Workshop on Health Security and Privacy. USENIX Association. Online at http://www.cs.dartmouth.edu/dfk/ papers/abstracts/mare-healthsec10.html.

MARE, S., SORBER, J., SHIN, M., CORNELIUS, C., AND KOTZ, D. 2011. Adaptive security and privacy for mHealth sensing. In *Proceedings of the USENIX Workshop on Health Security (HealthSec)*. Online at http://www.cs.dartmouth.edu/dfk/papers/mare-healthsec11.pdf.

MARTIN, L. 2008. Identity-based encryption and beyond. IEEE Security and Privacy 6, 62–64. Online at DOI 10.1109/MSP,2008.120.

MARY HITCHCOCK MEMORIAL HOSPITAL AND DARTMOUTH-HITCHCOCK CLINICS. 2009. The Dartmouth-Hitchcock Privacy Group policy statement on the privacy & confidentiality of patient information.

MATHUR, S., TRAPPE, W., MANDAYAM, N., YE, C., AND REZNIK, A. 2008. Radio-telepathy: extracting a secret key from an unauthenticated wireless channel. In *Proceedings of the ACM International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 128–139. DOI 10.1145/1409944.1409960.

MAYRHOFER, R. AND GELLERSEN, H. 2007. Shake well before use: Authentication based on accelerometer data. In *Proceedings of the International Conference on Pervasive Computing (Pervasive)*. Lecture Notes in Computer Science Series, vol. 4480. Springer-Verlag, 144–161. DOI 10.1007/978-3-540-72037-9_9.

MCDANIEL, P. AND RUBIN, A. 2000. A response to "Can we eliminate certificate revocation lists?". In *Proceedings of the International Conference on Financial Cryptography (FC)*, Y. Frankel, Ed. Lecture Notes in Computer Science Series, vol. 1962. Springer-Verlag, 245–258. DOI 10.1007/3-540-45472-1_17.

MERKLE, R. 1982. Method of providing digital signatures. US Patent 4309569. Online at http://patft.uspto.gov/netacgi/nph-Parser?patentnumber=4309569.

MESSMER, E. 2008. Health care organizations see cyberattacks as growing threat. Network World. Online at http://tinyurl.com/66b2py.

MF 2008. Common Framework for networked personal health information: Overview and principles. *Connecting for Health*. Online at http://connectingforhealth.org/phti/docs/Overview.pdf.

MFC 2009, Markle Foundation: Connecting for Health. Consumer consent to collections, uses, and disclosures of information. Online at http://connectingforhealth.org/phti/docs/CP3.pdf, visited Nov. 2009.

mH 2009, Wikipedia. mHealth. Online at http://en.wikipedia.org/wiki/Mhealth, visited Apr. 2009.

MHV 2008, Microsoft. The HealthVault web-based PHR. Online at http://www.healthvault.com, visited Nov. 2008.

MICALI, S. 2002. NOVOMODO: Scalable certificate validation and simplified PKI management. In *Proceedings of the PKI Research Workshop*. NIST. Online at http://www.cs.dartmouth.edu/_pki02/Micali/paper.pdf.

MID. 2009, Wikipedia. Mobile internet device. Online at http://en.wikipedia.org/wiki/Mobile Internet Device, visited May 2009.

MIŠIĆ, J. 2008. Enforcing patient privacy in healthcare WSNs using ECC implemented on 802.15.4 beacon enabled clusters. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE Computer Society Press, 686–691. DOI 10.1109/PERCOM.2008.28.

MOKBEL, M. F., CHOW, C.-Y., AND AREF, W. G. 2006. The new Casper: query processing for location services without compromising privacy. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*. VLDB Endowment, 763–774. Online at http://www.vldb.org/conf/2006/p763-mokbel.pdf.

MOLINA, A. D., SALAJEGHEH, M., AND FU, K. 2009. HICCUPS: Health information collaborative collection using privacy and security. In *Proceedings of the Workshop on Security and Privacy in Medical and Home-Care Systems (SPIMACS)*. ACM Press, 21–30. DOI 10.1145/1655084.1655089.

MONT, M. C., BRAMHALL, P., AND HARRISON, K. 2003. A flexible role-based secure messaging service: Exploiting IBE technology for privacy in health care. In *Proceedings of the International Workshop on Database and Expert Systems Applications*. IEEE Press, 432–437. DOI 10.1109/DEXA.2003.1232060.

MOORE, J. 2009. The feds and PHR privacy. Government Health IT. Online at http://www.govhealthit.com/ Articles/2009/01/26/The-feds-and-PHR-privacy.aspx.

MOTTA, G. H. AND FURUIE, S. S. 2003. A contextual role-based access control authorization model for electronic patient record. *IEEE Trans. Inf. Tech. Biomed. 7*, 3, 202–207. DOI 10.1109/TITB.2003.816562.

MPWG. 2009, Trusted Computing Group. Mobile Phone Work Group. Online at http://www. trustedcomputinggroup.org/developers/mobile, visited May 2009.

MTM. 2008, Trusted Computing Group. Mobile Phone Work Group Mobile Trusted Module Specification, Version 1.0. Online at http://www.trustedcomputinggroup.org/resources/mobile_ phone_workgroup_mobile_trusted_module_specification_version_10, visited June 2008.

MURALIDHAR, K. AND SARATHY, R. 2005. An enhanced data perturbation approach for small data sets. *Dec. Sci. 36*, 3, 513–529. DOI 10.1111/j.1540-5414.2005.00082.

NAHIT 2008. Defining key health information technology terms. Report to the Office of the National Coordinator for Health Information Technology. Online at http://www.nahit.org/images/pdfs/HITTermsFinalReport_051508.pdf.

NCVHS 2008. Individual control of sensitive health information accessible via NHIN. NCVHS letter to HHS Secretary. Online at http://www.ncvhs.hhs gov/080220lt.pdf.

NHS 2009a, UK National Health Service. Connecting for Health. Online at http://www.connectingforhealth.nhs.uk/, visited Mar. 2009.

NHS 2009b, UK National Health Service. Connecting for Health: Systems and services. Online at http://www.connectingforhealth.nhs.uk/systemsandservices, visited Mar. 2009.

NI, Q., LIN, D., BERTINO, E., AND LOBO, J. 2007a. Conditional privacy-aware role based access control. In *Proceedings of the European Symposium On Research In Computer Security (ESORICS)*. Lecture Notes in Computer Science Series, vol. 4734. Springer-Verlag, 72–89. DOI 10.1007/978-3-540-74835-9_6.

NI, Q., TROMBETTA, A., BERTINO, E., AND LOBO, J. 2007b. Privacy-aware role based access control. In *Proceedings of the ACM Symposium on Access Control Models and Technologies (SACMAT)*. ACM, 41–50. DOI 10.1145/1266840.1266848.

NISSENBAUM, H. 2004. Privacy as contextual integrity. Washington Law Review 79, 119–158. Online at http://www.nyu.edu/projects/nissenbaum/papers/washingtonlawreview.pdf.

NZHIPC. 2008. Health information privacy code 1994. New Zealand. 2008 revised edition. Online at http://www.privacy.org.nz/assets/Files/Codes-of-Practice-materials/HIPC-1994-2008-revised-edition.pdf.

NZPA. 1993. Privacy act 1993. New Zealand legislature, Public Act 1993 No. 28. Online at http://www.legislation.govt.nz/act/public/1993/0028/latest/096be8ed80604d98.pdf.

OECD. 1980. OECD guidelines on the protection of privacy and transborder flows of personal data. Online at http://preview.tinyurl.com/2of8ox.

ONC 2008. The nationwide privacy and security framework for electronic exchange of individually identifiable health information. Online at http://www.hhs.gov/healthit/privacy/framework.html.

OW 2009, Organized Wisdom. Organizedwisdom.com. Online at http://organizedwisdom.com, visited Oct. 2009.

PANG, J., GREENSTEIN, B., GUMMADI, R., SESHAN, S., AND WETHERALL, D. 2007. 802.11 user fingerprinting. In *Proceedings of the ACM International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 99–110. DOI 10.1145/1287853.1287866.

PARADISO, R., LORIGA, G., AND TACCINI, N. 2005. A wearable health care system based on knitted integrated sensors. *IEEE Trans. Inf. Tech. Biomed. 9*, 3, 337–344. DOI 10.1109/TITB.2005.854512.

PL 2008, Intel Research. PlaceLab project. Online at http://www.placelab.org/, visited Mar. 2008.

POUNDER, C. 2007. Why the APEC privacy framework is unlikely to protect privacy. Out-Law.com. Online at http://www.out-law.com/default.aspx?page=8550.

PRASAD, A. AND KOTZ, D. 2010. Can I access your data? Privacy management in mHealth. In *Proceedings of the USENIX Workshop on Health Security and Privacy*. USENIX Association. Online at http://www.cs.dartmouth.edu/~dfk/papers/abstracts/prasad-healthsec10.html.

PRASAD, A., SORBER, J., STABLEIN, T., ANTHONY, D., AND KOTZ, D. 2011. Exposing privacy concerns in mHealth. In *Proceedings of theUSENIX Workshop on Health Security (HealthSec)*. Online at http://www.cs.dartmouth.edu/~dfk/papers/prasad-healthsec11.pdf.

RAVICHANDRAN, R., BENISCH, M., KELLEY, P. G., AND SADEH, N. M. 2009. Capturing social networking privacy preferences. In *Proceedings of the International Symposium on Privacy Enhancing Technologies (PETS)*. Lecture Notes in Computer Science Series, vol. 5672. Springer-Verlag, 1–18. DOI 10.1007/978-3-642-03168-7_1.

RIEDL, B., NEUBAUER, T., GOLUCH, G., BOEHM, O., REINAUER, G., AND KRUMBOECK, A. 2007. A secure architecture for the pseudonymization of medical data. In *Proceedings of the International Conference on Availability, Reliability and Security (ARES)*. IEEE press, 318–324. DOI 10.1109/ARES.2007.22.

RIVEST, R. L. 1998. Can we eliminate certificate revocations lists? In *Proceedings of the International Conference on Financial Cryptography (FC)*, R. Hirschfeld, Ed. Lecture Notes in Computer Science Series, vol. 1465. Springer-Verlag, 178–183. DOI 10.1007/BFb0055482.

ROUSE, W. B. 2008. Health care as a complex adaptive system: Implications for design and management. *The Bridge 38*, 1. Online at http://www.nae.edu/nae/bridgecom.nsf/weblinks/MKEZ-7CLKRV? OpenDocument.

SAFE. 2010. U.S. Department of Commerce. Welcome to the U.S.-EU & Swiss safe harbor frameworks. Online at http://www.export.gov/safeharbor, visited Oct. 2010.

SAHAI, A. AND WATERS, B. 2005. Fuzzy identity-based encryption. In *Proceedings of Advances in Cryptology (EUROCRYPT)*. Lecture Notes in Computer Science Series, vol. 3494. Springer-Verlag, 457–473. DOI 10.1007/11426639 27.

SALTZER, J. H. AND SCHROEDER, M. D. 1975. The protection of information in computer systems. *Proc. IEEE 63*, 9, 1278–1308. DOI 10.1109/PROC,1975,9939.

SAMARATI, P. 2001. Protecting respondents' identities in microdata release. *IEEE Trans. Knowl. Data Eng. 13*, 6, 1010–1027. DOI 10.1109/69.971193.

SANDHU, R. S., COYNE, E. J., FEINSTEIN, H. L., AND YOUMAN, C. E. 1996. Role-based access control models. *IEEE Comput. 29*, 2, 38–47. DOI 10.1109/2.485845.

SANKAR, P. AND JONES, N. L. 2005. To tell or not to tell: primary care patients' disclosure deliberations. *Arch. Intern. Med. 165*, 20, 2378–2383. DOI 10.1001/archinte.165.20.2378.

SCHOLL, M., STINE, K., HASH, J., BOWEN, P., JOHNSON, A., SMITH, C. D., AND STEINBERG, D. I. 2008. An introductory resource guide for implementing the Health Insurance Portability and Accountability Act (HIPAA) security rule. Tech. Rep. 800-66-Rev1, National Institute of Standards and Technology. Oct. Online at http://csrc.nist.gov/publications/nistpubs/800-66-Rev1/SP-800-66-Revision1.pdf.

SCHWINGENSCHLÖGL, C., EICHLER, S., AND MÜLLER-RATHGEBER, B. 2006. Performance of PKI-based security mechanisms in mobile ad hoc networks. *Int. J. Electron. Commun. 60*, 1, 20–24. DOI 10.1016/j.aeue.2005.10.004.

SH 2008, University of Rochester. Smart Home project at Center for Future Health. Online at http://www.futurehealth.rochester.edu/smart_home, visited Mar. 2008.

SINCLAIR, S. AND SMITH, S. W. 2008. Preventative directions for insider threat mitigation via access control. In *Insider Attack and Cyber Security: Beyond the Hacker*. Advances in Information Security Series, vol. 39. Springer-Verlag, 173–202. DOI 10.1007/978-0-387-77322-3_10.

SINGELÉE, D. AND PRENEEL, B. 2006. Location privacy in wireless personal area networks. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*. ACM, 11–18. DOI 10.1145/1161289.1161292.

SOLWORTH, J. A. 2008. Instant revocation. In *Public Key Infrastructure*. Lecture Notes in Computer Science Series, vol. 5057. Springer-Verlag, 31–48. DOI 10.1007/978-3-540-69485-4_3.

SRINIVASAN, V., STANKOVIC, J., AND WHITEHOUSE, K. 2008. Protecting your daily in-home activity information from a wireless snooping attack. In *Proceedings of the Conference on Ubiquitous Computing (UbiComp)*. ACM, 202–211. DOI 10.1145/1409635.1409663.

SRINIVASAN, V., STANKOVIC, J., AND WHITEHOUSE, K. 2010. Using height sensors for biometric identification in multi-resident homes. In *Proceedings of the International Conference on Pervasive Computing (Pervasive)*. Lecture Notes in Computer Science Series, vol. 6030. Springer, Berlin Heidelberg, 337–354. DOI 10.1007/978-3-642-12654-3_20.

SRIRAM, J., SHIN, M., CHOUDHURY, T., AND KOTZ, D. 2009a. Activity-aware ECG-based patient authentication for remote health monitoring. In *Proceedings of the International Conference on Multimodal Interfaces and Workshop on Machine Learning for Multi-modal Interaction (ICMI-MLMI)*. ACM, 297–304. DOI 10.1145/1647314.1647378.

SRIRAM, J., SHIN, M., KOTZ, D., RAJAN, A., SASTRY, M., AND YARVIS, M. 2009b. Challenges in data quality assurance in pervasive health monitoring systems. In *Future of Trust in Computing*, D. Gawrock, H. Reimer, A.-R. Sadeghi, and C. Vishik, Eds. Vieweg+Teubner Verlag, 129–142. DOI 10.1007/978-3-8348-9324-6 14.

STANFORD, V. 2002. Pervasive health care applications face tough security challenges. *IEEE Pervas. Comput. 1*, 2, 8–12. DOI 10.1109/MPRV.2002.1012332.

STEINBROOK, R. 2009. Health care and the American Recovery and Reinvestment Act. *New Eng. J. Med. 360*, 11, 1057–1060. DOI 10.1056/NEJMp0900665.

SUN, Y., LA PORTA, T. F., AND KERMANI, P. 2009. A flexible privacy-enhanced location-based services system framework and practice. *IEEE Trans. Mobile Comput. 8*, 3, 304–321. DOI 10.1109/TMC.2008.112.

SUNDARAM, B. AND CHAPMAN, B. 2005. A grid authentication system with revocation guarantees. In *Proceedings of the Symposium on High Performance Computing (HiPC)*. Lecture Notes in Computer Science Series, vol. 3769. Springer, 508–517. DOI 10.1007/11602569_52.

SWEENEY, L. 2002. *k*-anonymity: A model for protecting privacy. *Int. J. Uncert., Fuzz., Knowl.-Based Syst. 10*, 5, 557–570. DOI 10.1142/S0218488502001648.

TAN, C. C., WANG, H., ZHONG, S., AND LI, Q. 2009. IBE-lite: A lightweight identity-based cryptography for body sensor networks. *IEEE Trans. Inf. Tech. Biomed. 13*, 6, 926–932. DOI 10.1109/TITB.2009.2033055.

TPM. 2009, Trusted Computing Group (TCG). Trusted Platform Module. Online at http://www.trustedcomputinggroup.org/developers/trusted_platform_module, visited May 2009.

UNIVERSITY OF WASHINGTON. 2008. Assisted Cognition project. http://www.cs.washington.edu/Assistcog. (last accessed 3/08).

VADEHRA, S. 2011, Kan & Krishme, Attorneys at Law. India: Data protection and the IT Act India. Online at http://www.gala-marketlaw.com/joomla4/index.php?option=com_content&view=article&id=261&Itemid=138, visited Jan. 2011.

VARSHAVSKY, A., LAMARCA, A., AND DE LARA, E. 2007a. Enabling secure and spontaneous communication between mobile devices using common radio environment. In *Proceedings of the Workshop on Mobile Computing Systems and Applications (HotMobile)*. ACM, 9–13. DOI 10.1109/HotMobile.2007.12.

VARSHAVSKY, A., SCANNELL, A., LAMARCA, A., AND DE LARA, E. 2007b. Amigo: Proximity-based authentication of mobile devices. In *Proceedings of Ubiquitous Computing (UbiComp)*. Lecture Notes in Computer Science Series, vol. 4717. Springer-Verlag, 253–270. DOI 10.1007/978-3-540-74853-3_15.

VARSHNEY, U. 2007. Pervasive healthcare and wireless health monitoring. *Mobile Netw. Appl. 12*, 2-3, 113–127. DOI 10.1007/s11036-007-0017-1.

VITALETTI, A. AND PALOMBIZIO, G. 2007. Rijndael for sensor networks: Is speed the main issue? *Electron. Notes Theoret. Comput. Sci. (ENTCS) 171*, 1, 71–81. DOI 10.1016/j.entcs.2006.11.010.

WANG, Q., SHIN, W., LIU, X., ZENG, Z., OH, C., ALSHEBLI, B. K., CACCAMO, M., GUNTER, C. A., GUNTER, E., HOU, J., KARAHALIOS, K., AND SHA, L. 2006. I-Living: An open system architecture for assisted living. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC)*. Vol. 5. IEEE press, 4268–4275. DOI 10.1109/ICSMC.2006.384805.

WANG, W., MOTANI, M., AND SRINIVASAN, V. 2008. Dependent link padding algorithms for low latency anonymity systems. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM, 323–332. DOI 10.1145/1455770.1455812.

WATRO, R., KONG, D., CUTI, S.-F., GARDINER, C., LYNN, C., AND KRUUS, P. 2004. TinyPK: securing sensor networks with public key technology. In *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*. ACM, 59–64. DOI 10.1145/1029102.1029113.

WEERASINGHE, D., ELMUFTI, K., RAJARAJAN, M., AND RAKOCEVIC, V. 2007. Securing electronic health records with novel mobile encryption schemes. *Int. J. Electron. Healthcare 3*, 4, 395–416. DOI 10.1504/IJEH.2007.015320.

WONG, F.-L. AND STAJANO, F. 2005. Location privacy in Bluetooth. In *Proceedings of the European Workshop on Security and Privacy in Ad-hoc and Sensor Networks (ESAS)*. Lecture Notes in Computer Science Series, vol. 3813. Springer-Verlag, 176–188. DOI 10.1007/11601494_15.

WRIGHT, C. V., BALLARD, L., COULL, S. E., MONROSE, F., AND MASSON, G. M. 2010. Uncovering spoken phrases in encrypted voice over IP conversations. *ACM Trans. Inf. Syst. Sec. (TISSEC) 13*, 4, 35:1–35:30. DOI 10.1145/1880022.1880029.

WRIGHT, C. V., COULL, S. E., AND MONROSE, F. 2009. Traffic morphing: An efficient defense against statistical traffic analysis. In *Proceedings of the Annual Symposium on Network and Distributed System Security (NDSS)*. Internet Society. Online at http://www.isoc.org/isoc/conferences/ndss/09/pdf/14.pdf.

XIAO, Y., RAYI, V. K., SUN, B., DU, X., HU, F., AND GALLOWAY, M. 2007. A survey of key management schemes in wireless sensor networks. Computer Communications 30, 11-12, 2314–2341. Special issue on security on wireless ad hoc and sensor networks, DOI 10.1016/j.comcom.2007.04.009.