Privacy Challenges in Patient-Centric Health Information Systems

Anupam Datta[1], Nipun Dave[3], John Mitchell[3], Helen Nissenbaum[2], Divya Sharma[1]

1:Carnegie Mellon University
{danupam, divyasharma}@cmu.edu

2:New York University
hfn1@nyu.edu

3:Stanford University
{nipund, John.Mitchell}@stanford.edu

***Abstract*—Patient Health Record (PHR) systems offer great promise but raise significant philosophical, cultural, legal, and technical challenges. In hopes of furthering debate on key issues, we explain some central questions about the role, purpose, and policies associated with these systems. We also propose a framework for addressing policy questions and candidate technology that we believe may sharpen policy discussion and allow PHR systems to adhere to policies they adopt.**

***Keywords*-Privacy, policy, contextual integrity, conformance**

I. Introduction

The emergence of patient-centric health information systems, including Personal Health Record (PHR)[1] sites such as Google Health [1] and Microsoft HealthVault [2], holds great promise for empowering patients and ensuring more effective delivery of health care. At the same time, these systems raise significant patient privacy challenges because organizations running successful PHRs will have access to sizable databases of personal health information. This aggregate health information has economic value to insurance companies, pharmaceuticals, and others, creating economic incentives for flows of personal health information that may not align with patients' interests. While health care providers, such as hospitals and clinics, are regulated by HIPAA [8], there is no comparable comprehensive regulation that meaningfully constrains transmission and use of personal health information by PHRs or related patient-centric health information systems.

In order for PHR sites to serve patients and the healthcare system that serves them, there is a pressing need to develop *socially acceptable privacy policies* that govern the flow and use of personal health information in patient-centric health information systems, along with *technology* and legal policy to support the enforcement of such policies. This position paper recognizes this need, raises topics for debate, and outlines an approach for addressing healthcare privacy requirements through a combination of social and technology-based methods, building on prior work by some of the authors. Our privacy position for patient-centric health information systems is founded on *contextual integrity* [7], a philosophical theory of privacy that goes beyond the commonly held position that privacy is about control over one's information. The prescriptive component of contextual integrity provides a framework for determining what kind of privacy policies may be socially acceptable in this context, which we elaborate on in Section II. A specific technology problem that this work addresses is that of representation of policies for PHRs in a policy language with precise semantics and automated techniques for enforcement of such policies [4], [5]. The policy language, described in Section III formalizes key concepts from the descriptive component of contextual integrity.

[1]A PHR is "an electronic application through which individuals can access, manage, and share their health information, in a private, secure, and confidential environment; personal data created, developed, maintained, and/or provided by individuals about themselves" [9].

II. Privacy Position

A Personal Health Record is generally a health record that is initiated and maintained by an individual. Google Health, for example, claims to store information "securely and privately" and let patients "always control how it's used" [1]. Microsoft HealthVault similarly proposes to allow individuals to "take charge" and "make more informed health decisions" [2]. Both sites promise to help individuals gather and organize medical records. Patients may naturally expect that this collected information will help them understand their health issues more clearly, and also allow them to provide information about past diagnosis and treatment with medical professionals; Google explicitly highlights sharing information with "doctors or caregivers" [1].

One basic issue is the degree to which an individual may restrict visibility into information they store in their PHR. In the commonly held view of privacy as a right to control information about oneself, adopted by privacy advocates including Deborah Peel of Patients Privacy Rights, patient control would seem to effectively address privacy concerns. However, it is not clear how complete control could be achieved, it is not clear that current sites promise it, and it is debatable whether complete control is actually in the best interests of individuals or the public good. Certainly no individual wishes to be asked directly whenever someone wishes to access their health record. Further, when aggregate statistics are calculated, there is room for debate as to whether release of those statistics constitute use of personal health data. With regard to individual control over their health information, Google's privacy policy allows use in other Google Products; although data will not be used

to customize ads, there are apparently no further explicit restriction on the cross-product use of data. Finally, epidemics and spread of certain diseases are currently tracked my government health agencies, and it is likely that laws requiring notification or tracking of certain diseases will be applied to PHRs, in the interest of the public good. We therefore question the simple view that equates privacy with individual control. Instead, we propose evaluation and debate regarding a broader view of privacy based on the theory of contextual integrity [7].

According to the framework of contextual integrity, control is an important aspect of privacy, but only partially so. The core claim by contrast, is that privacy protection amounts to protection of appropriate information flow. Appropriateness, here, is modeled by context-relative information norms, or rules, that prescribe the flow of information from one party to another based on the capacities in which the parties are acting (parties being sender, recipients, and subjects of the information), the type of information, and the constraints under which the information flows. In some cases, control by the subject may be an appropriate constraint on the flow of information, but in other cases it may not be. The norms, or rules, that a society supports through mechanisms like laws — though there may be others — should be those that promote the values, ends, and purposes of the background context. We therefore take the following privacy position:

Privacy is not just about patients' control over their information. It is better characterized in terms of norms that govern the flow of personal health information in a manner that promotes the values, ends, and purposes of the health care context.

Accordingly, in determining privacy rules for PHRs, we must understand the ways PHRs function in healthcare, what ends and purposes they serve, and what values they promote (or impede). In short, we propose a public debate to derive suitable privacy rules for PHRs based on the ways PHRs function in the larger healthcare environment. An important step along the way will be to articulate, with care, the function of PHRs. Because the PHR is a relatively new concept, we anticipate that this will be part of a lively debate.

III. Technology for Policy Enforcement

Technology can play an important role in ensuring that privacy policies associated with PHRs and other patient-centric health information systems are expressed precisely and unambiguously, as well as in ensuring that processes and people in such organizations act in a manner that is compliant with the stated policies.

Our prior work on formal expression of HIPAA and its enforcement in organizational processes in hospitals provides a useful starting point for this effort [4], [5]. Specifically, we formalized key concepts of contextual integrity to develop a model of organizational processes. Using a model of *actions* that transmit personal information from a sender in one role to a receiver in a possibly different role, agents may accumulate and send different types of personal information they receive. These messages represent emails, web forms, database entries, workflow data structures, and arguments to actions. Since agents may act independently, with different motives, we express privacy and utility (contextual purpose) goals using a form of alternating-time temporal logic, which we call the *Logic of Privacy and Utility (LPU)*, interpreted over the concurrent game structure [3] of agent actions. In this logical setting, privacy is a trace property expressible in LTL [6], while utility requires that agents have strategies to achieve certain useful outcomes, and is therefore expressed naturally using the stronger ATL* [3] path quantifiers. We also formulate processes in temporal logic, by associating a *responsibility* to each agent role. Within this setting, policy enforcement is achieved using a combination of design-time analysis based on model-checking and posthoc auditing for detection of policy violation and blame assignment to principals whose irresponsible actions caused the violation.

Policy enforcement for patient-centric health information systems promises to require new techniques to address additional challenges associated with flows of aggregate personal health information and the incentives that drive such flows in the complex health information system context.

References

[1] Google health. [Online]. Available: http://www.google.com/intl/en-US/health/about/index.html

[2] Microsoft healthvault. [Online]. Available: http://www.healthvault.com/Personal/index.html

[3] R. Alur, T. A. Henzinger, and O. Kupferman, "Alternating-time temporal logic," *J. ACM*, vol. 49, no. 5, pp. 672–713, 2002.

[4] A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum, "Privacy and contextual integrity: Framework and applications," in *SP '06: Proceedings of the 2006 IEEE Symposium on Security and Privacy*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 184–198.

[5] A. Barth, J. C. Mitchell, A. Datta, and S. Sundaram, "Privacy and utility in business processes," in *CSF*, 2007, pp. 279–294.

[6] Z. Manna and A. Pnueli, *Temporal Verification of Reactive Systems: Safety*. Springer–Verlag, 1995.

[7] H. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2010.

[8] Office for Civil Rights, "Summary of the HIPAA privacy rule," US Department of Health & Human Services, 2003.

[9] C. Safran, M. Bloomrosen, W. E. Hammond, S. Labkoff, S. Markel-Fox, P. C. Tang, D. E. Detmer, and P. Expert, "Toward a national framework for the secondary use of health data: an american medical informatics association white paper," *Journal of the American Medical Informatics Association*, vol. 14, pp. 1–9, October 2006.