usenix ;login:
DECEMBER 2012 VOL. 37, NO. 6
ELECTRONIC SUPPLEMENT

Conference Reports

3rd USENIX Workshop on Health Security and Privacy (HealthSec '12)

Bellevue, WA
August 6-7, 2012

Keynote Address

Summarized by Daniel Yang Li (dyli@uw.edu)

Of Codes, Genomes, and Electronic Health Records: It's Only Sensitive If It Hurts When You Touch It

Dan Masys, University of Washington

Professor Daniel Masys, from the Department of Biomedical Informatics and Medical Education of the University of Washington, shared his experience in building secure systems for health care. It is noteworthy that Professor Masys was a pioneer in building systems for health care, dating back to the 1990s.

Masys first listed the topics he wanted to talk about. In particular, he pointed out that the systems for health care are built on trust over health data. He mentioned the HIPAA/HITECH rules, and emphasized the effects and spirit of such rules. Masys said that the security rules and the privacy rules are distinct.

Later, Masys discussed the security vulnerabilities in health care. A famous example is that by using data mining techniques, it is possible to re-identify patients. Masys also presented the security issues related to genome sequences. It is well-known that some parts of a genome sequence may correlate with a particular disease. On the other hand, the relationship between genome sequences and diseases is not fully understood, so using genome sequences to infer a certain disease is difficult in reality.

Masys examined the properties of health care data systems and claims that designing such systems is hard. It is inappropriate to ask questions such as who owns the data, because the ownership issue is so complicated. In the end, Masys restated that it is crucial to take security and privacy into account when designing health care systems and that it is the job of the security community to design better systems.

There were questions about the complexity of checking the integrity of health data; hacking medical reports; cloud security for genome data; and big data. Masys answered that he did not have specific answers to these questions. Rather, he called on the security community to attack these problems in a joint effort. Masys said that many research papers could be produced by working on the many unsolved security issues in health data management, and it is the duty of the whole community to solve them.

Privacy

Summarized by Tamara Denning (tdenning@cs.washington.edu)

Neuroimage Data Sets: Rethinking Privacy Policies

Nakeisha Schimke and John Hale, Institute of Bioinformatics and Computational Biology, University of Tulsa

Nakeisha Schimke described the current state of data sharing and deidentification/reidentification to the field of neuroimaging, then calls for a more unified, principled approach in the community. There are many advantages to sharing neuroimaging data—including the fact that there is a limited amount of patient and study data available—but privacy concerns on the part of participants and other stakeholders complicates the space. While metadata can be filtered before sharing data, there is no uniform approach to deidentifying the visual data (skull stripping, quickshear, and MRI defacer are some individual approaches).

Furthermore, image alteration interferes with the quality of images and the accuracy of a study, so the community should perform additional studies to understand the privacy-threat landscape better and perform the minimum necessary alterations to images; no good baseline study exists dealing with the reidentifiability of neuroimaging data. Such a study should be performed to inform a unified privacy policy dealing with neuroimaging.

An audience member asked which organization can or should create and enforce such privacy policies dealing with neuroimaging data. Schimke suggested the NIH as an appropriate organization.

Protecting Web-Based Patient Portal for the Security and Privacy of Electronic Medical Records

Xiaowei Li and Yuan Xue, Vanderbilt University

Xiaowei Li presented and described OpenEMR, an open-source patient portal for electronic medical records, with particular emphasis on the privacy and security of EMRs. OpenEMR employs a two-tier security model: Request Blocker (based on BLOCK) inspects Web requests to look for suspicious behavior, and EMR Protector (based on SENTINEL) isolates the EMR from the patient portal. Both components consult the central decision engine, which implements specified security policies. Blocker and Protector have been implemented for general-purpose Web applications.

The proposed approach is that security specifications can be learned from the patterns of Web access; for example, learned policies could include rule-based specification (e.g., professional role in medicine) and evidence-based specification (e.g., a heart attack patient's record is accessed by people from many different departments).

One audience member asked how OpenEMR differs from normal Web application security. Li answered that the security specifications for patient portals need to reflect clinical access policies, which can be abnormal or complicated. Another question was whether or not this system addresses side-channel attacks; it does not.

Vis-à-vis Cryptography: Private and Trustworthy In-Person Certifications

Ian M. Miers and Matthew Green, Johns Hopkins University; Christoph U. Lehmann, MD, Johns Hopkins University School of Medicine; Aviel D. Rubin, Johns Hopkins University

Ian Miers presented a paper that investigates a cryptographic protocol for presenting in-person certifications that are deniable and linked to an identity. The particular problem space addressed in this paper was the challenge of proving to another individual via an in-person cell phone app that one has been tested for STIs. This is a space where certifiability from a testing clinic is desirable, since both faking false positives and false negatives can negatively impact other individuals. Additionally, deniability is desirable, since the act of showing someone else one's status should not be used after-the-fact as evidence of an interaction. Furthermore, the certifiability should be linked with a picture, so that the results shown are guaranteed to belong to the individual in question.

The researchers are also investigating extensions to this protocol, such as securely and privately being able to inform individuals after the fact if they were exposed to an STI. Other potential extensions include marking an individual as exposed (if they were exposed and have not/cannot yet be definitively tested) or to display an individual's risk factor.

Someone asked whether the app is being tested and deployed. Miers answered that they are looking for a small college that would be willing to serve as a testbed for this app, since it would work best if there is a critical mass of users.

Panel: Privacy-Protecting Management of Health Data

Summarized by Daniel Yang Li (dyli@uw.edu)

Privacy-Protecting Management of Health Data

Xiaofeng Wang, Indiana University; Darren Larcey, Johns Hopkins University; Haixu Tang, Indiana University, and Xiaoqian Jiang, UC San Diego

Wang talked about the future challenge of health data management and mentioned the genome privacy problem. Wang emphasized the importance of protecting privacy around genetic records, and of designing privacy-preserving genome analysis. Larcey made points from a practical perspective on health data management, speaking of the sources of clinical research, emerging security controls, and directions for further research.

Tang, an expert in bioinformatics, expressed concern about health data management based on his research. Integrated personal omics (the study of a body of information) profiling needs to incorporate a better understanding of health data management. Jiang emphasized the balance between utility and privacy. The idea of differential privacy addresses the tradeoff between utility and privacy is.

Carl Gunter (University of Illinois) asked two closely related questions on the privacy budgets in differential privacy and in multiparty cases. Jiang admitted that these are open problems in the field and, so far, there are no satisfying answers.

Biometrics

Summarized by Tamara Bonaci (tbonaci@uw.edu)

Who Wears Me? Bioimpedance as a Passive Biometric

Cory Cornelius, Jacob Sorber, Ronald Peterson, Joe Skinner, Ryan Halter, and David Kotz, Institute for Security, Technology, and Society, Dartmouth College

Cory Cornelius presented a new sensor for passive identification of persons wearing a health monitoring system. This work is motivated by the requirement that commonly used mobile and wearable health monitoring systems need to identify the wearer correctly in order to label and store the collected data properly. While patient identification systems already exist, these systems can easily be tricked into incorrectly identifying a wearer without the patient's active engagement. Thus, there is a need to develop a passive biometric sensor, which will measure features universally available yet unique to each person (such as a fingerprint). Additionally, these features should be stable over time, easy to measure, but difficult to forge.

One such feature is bioimpedance, a measure of how the body tissue opposes an applied alternating current of small amplitude. It is a physiological property related to a tissue's resistance to electrical current flow and its ability to store electrical charge, and it is unique to every person.

Cornelius et al. developed a passive biometric device that identifies a wearer based on the measured bioimpedance. In this talk, Cornelius discussed two methods of measuring bioimpedance. He then presented a prototype bioimpedance sensor and analyzed experimental results, proving the feasibility of the proposed method.

How was the cohort of 46 human subjects chosen for experiments, and were the experiments repeated? The subjects were chosen randomly, and the experiments were indeed repeated. What was their experience with the institutional review board (IRB)? It was positive; getting the IRB's approval required only a minimal effort, since they already had a similar study approved. Could the presented experimental results be considered biased, given that the tests were

only performed on students? With respect to their bioimpedance, the group of students might be more homogeneous than the members of the same household, thus making the recognition process only harder. What classifiers were used for identification, and what are the next steps of the project? They used linear and Bayesian classifiers but other classifiers are certainly possible. The next step is the development of a wearable sensor, as well as feature engineering.

Body Area Network Security: Robust Key Establishment Using Human Body Channel

Sang-Yoon Chang, Yih-Chun Hu, Hans Anderson, and Ting Fu, University of Illinois at Urbana-Champaign; Evelyn Y. L. Huang, University of Illinois at Chicago

Sang-Yoon Chang presented novel results in a reliable and secure exchange of cryptographic keys in body area networks (BANs). A BAN consists of several sensors in physical contact with the human body. The sensors, in addition to measuring a person's physiological parameters, periodically exchange data between them. The wireless connection between the sensors introduces vulnerabilities, however, such as eavesdropping and impersonation. To mitigate potential attacks on BANs, therefore, sensors establish secure communication by exchanging a cryptographic key.

The existing schemes for secure and reliable secret exchange in BANs typically rely on the randomness of physiological values. In this work, Chang et al. examined the process of measuring physiological values, and analyzed the practicality of using those measurements for secret sharing in BAN applications. Chang showed that the ECG signals, the physiological signals most commonly used for BAN security, are highly dependent on the sensors' deployment location, and exhibit noisy measurements on muscular body regions. Moreover, he showed how the existing secret exchange methods can leak private information to an outside attacker, who can remotely measure ECG signals. Chang then presented a novel secret exchange method for BANs, which uses voltage signals operating below the human body action potential. He analyzed the feasibility of the proposed method by discussing several experimental results, and presented the Shannon capacity of the proposed method, which was computed combining the model of the human body channel and empirical data.

How would the proposed scheme be used to exchange cryptographic secrets of 80 bits or more, given that the Shannon capacity of the proposed method is 11 bits/day? Chang acknowledged that the computed capacity is not high, but pointed out it is only a lower bound, and that the authors have used a very conservative approach in deriving that bound. The proposed bound is not too restrictive, since secret updates will only be needed occasionally, and if that is the case, then secrets can be updated during a patient's stay in the hospital.

Can the proposed method be enhanced to increase the Shannon capacity, for example by augmenting a measured physiological signal with a random voltage signal? Chang replied that might be possible, but the authors have not investigated it yet. He pointed out, however, that the proposed method considers a stronger adversary, who can remotely measure physiological signals. Would the proposed method fail if the attacker had a contact with the skin—for example, if he would use a door knob as a point of contact? Chang acknowledged the secret exchange would indeed be compromised in that case, but that this case falls outside the considered adversarial model.

Keynote

Summarized by Sang-Yoon Chang (chang6@illinois.edu)

Security Lessons Learned from HIPAA Enforcement

Adam H. Greene, Davis Wright Tremaine LLP

Adam Greene discussed how information technology gets enforced on the legislative level and in practice. Greene, who used to work at the US Department of Health and Human Services (or HHS, the organization that enforces HIPAA), provided various data analyses based on legal complaints and security-related events to further analyze enforcement trends.

Greene began the keynote address with background information about HIPAA, such as the transition between organizations and the shift in enforcement focus as a result of patients' demands and concerns. HIPAA was published in 2003 with a compliance date of April 2005. It was initially enforced by the HHS Center for Medicare and Medicaid services (CMS) but then was taken over by HHS Office for Civil Rights (OCR) in July 2009.

Analyzing the security reports from year 2005 to year 2011, Greene found that there were not many security complaints but a number of privacy complaints, because the notion of security is less transparent to patients than privacy. The patients also seemed to be more concerned about neighbors' mishaps and less intrigued by how the patient information gets mapped into the hospital system. The top security issues of OCR also reflected these patients' concerns.

There has been a surge of breach reports between September 2009 and June 2012. From the analysis of these reports, Greene presented three trends. First, the breach data implies that the vulnerability revolves more around hospital employees (in the form of theft and unauthorized access/disclosure/loss) than hackers. Second, paper accounts for the highest number of breaches (paper documents subject to

theft or improper disposal or misdirected faxes); however, backup tapes are the biggest source location of breaches in terms of number of individuals affected, and large-scale breaches are accomplished through digital sources rather than paper. Third, many large breaches are caused by business associates rather than the covered entity. Due to the need for a risk analysis of business associates, there have been substantial privacy and security audits that are more proactive in nature (rather than incident driven). Greene also mentioned that these reviews include site visits and audit reports, and the subjects are extended to organizations that contribute to health plans, health care providers, and health care clearinghouses. The audit result is that security tends to be a big issue (consisting of 65% of findings) for adverse finding and the sources of such a covered entity are predominantly health care providers (81%).

After listing some real-life incidents that lead to HHS settlement, penalties, and criminal cases, Greene discussed the necessity to encrypt all health data, to focus on large data sets (including backup tapes and spreadsheets), and to pay closer attention to VIPs and sensitive information. He concluded by noting the HHS tendency to focus on systematic problems and its past use of voluntary enforcement, although settlements are increasing.

One person asked about the deficiency of access management and Greene agreed that there should be an improvement in access management. A couple of questions led to a discussion about whether these trends are also applicable to other nations. According to Greene, some nations (European and Israel) have data digitally stored and have a centralized system, which shifts the nature of the breach cause.

System Design

Summarized by Tamara Bonaci (tbonaci@uw.edu)

Information Security of Patient-Centered Services Utilizing the German Nationwide Health Information Technology Infrastructure

Tobias Dehling and Ali Sunyaev, University of Cologne, Germany

Ali Sunyaev presented the German Nationwide Health Information Technology Infrastructure (HTI), currently under development. The HTI is a health information technology, expected to enhance the existing German health care system and to provide secure patient-centered services. Patient-centered services are a special kind of an information technology, designed to manage personal medical information (PMI) and to offer advanced and personalized services and information to patients, based on the available PMI.

Due to sensitivity of personal medical information, but also the severity of possible consequences if the PMI data is breached, patient-centered services will need to implement several levels of security and privacy measures. Sunyaev analyzed the existing HTI system and compared it with the requirements for secure patient-centered services. He proposed the key management processes implemented by the HTI can simply be extended to patient-centered services in order to ensure confidentiality of stored and exchanged data. He further proposed that to maintain confidentiality of information, patient-centered services should not use true identities of users and should enforce the minimization principle and utilize the available information only in secure environments. He also proposed employing appropriate backup strategies, as well as encryption and signatures, in order to improve information integrity.

How will access control in patient-centered services work, since a large number of people can be expected to have access to a patient's PMI data? Specifically, how can potential abuse of patient privacy be prevented? Sunyaev responded that the HTI is already implemented such that nobody can access a patient's data without the patient' consent. Moreover, the HTI is a two-card access system where it is assumed the patient gives consent by presenting his/her unique health-care card, eHC, to a health service provider. What happens if a patient forgets to bring the eHC, or in the case of emergency? Such situations are regulated by rules, and Ali said part of the eHC was intentionally left unencrypted in case of emergency. And in Germany, in the case of emergency, a patient's health card is not required in order to treat the patient.

An Analysis of HIPAA Breach Data

Patrick Morrison and Laurie Williams, North Carolina State University

The Health Insurance Portability and Accountability Act (HIPAA) is a regulation governing the protection of personal health information (PHI) in the United States, and defining measures when institutions fail to protect the data. As part of the HITECH Act of 2009, organizations experiencing exposure of PHI are required to report to the Department of Health and Human Services all PHI-related breaches, and if a breach affects 500 or more individuals, a breach report is publicly posted on an Office of Civil Rights Web page.

Patrick Morrison presented the analysis of 392 reports on breaches that occurred between September 2009 and November 2011. The focus of the analysis was on the type of breach (theft, hacking/IT incident or loss, etc.) and the location of the breach (laptop, paper, network, etc.). Morrison noted that every location category was subject to a theft incident, but the easier it was to move a device or media containing PHI data, the greater was the probability of a breach, regardless of the type of breach. He further reported that the most commonly breached media type was paper, and that in

almost all reported breaches no encryption method was used. Morrison then proposed several strategic responses to potential PHI breaches, which included disabling print functionalities, encrypting PHI data at the time of its creation, and limiting the amount of PHI data stored on portable devices.

Could cloud services limit the needed number of copies of patient records? Morrison acknowledged that was an interesting proposal, but pointed out any attempt to use clouds as PHI data storage and management facilities should be carefully examined, due to known open issues with cloud services' privacy. Could he comment on the proposal to print encrypted data such that only a person wearing so-called "decoder" glasses could read it, the printed text appearing illegible to everyone else? He thought the proposed idea was quite interesting, but should be explored further.

Security Risks, Low-Tech User Interfaces, and Implantable Medical Devices: A Case Study with Insulin Pump Infusion Systems

Nathanael Paul, University of Tennessee, Oak Ridge National Laboratory; Tadayoshi Kohno, University of Washington

Portable implantable medical devices are increasingly being used in modern health care. Recently, a set of potential security issues exploiting wireless control interface of these devices has been identified. In this work, Nathanael Paul presented a novel class of security threats, targeting a low-tech user interface (UI) of implantable medical devices. Paul identified the following characteristics as distinguishing points between the threats targeting wireless control interface and those targeting user interface: delayed effect on patients, the fact that UI-attacks do not require technical sophistication, and the fact that UI-attacks can be targeted.

Paul analyzed an example UI interface of an insulin pump infusion system, where the pump device in a "patch pump" architecture is typically programmed using a remote wireless device similar to a smartphone. The remote control stores system settings and programs for insulin delivery, and transmits commands to the patient's pump device. He reported 20 potential settings of the remote devices that can negatively affect a patient. For example, a malicious entity with access to the remote control could easily change the insulin level of the pump, or the amount of insulin per unit of time, referred to as a basal rate.

Paul suggested that work mitigating the UI-attacks should focus on both prevention and detection of possible exploits of the user interface. In prevention, a better authentication might stop unwanted changes from occurring, and as one possible way of improving authentication, Paul suggested augmenting insulin pump components by adding context to system events. In order to detect attacks, Paul proposed improvements in system event recording.

Did Paul believe the identified problems are UI-design-related or security-related? He responded that while some of the identified issues are indeed UI-design issues, they are also important safety and security issues, and should be addressed accordingly.

Access Control

Summarized by Tamara Denning (tdenning@cs.washington.edu)

Development of a System Framework for Implementation of an Enhanced Role-Based Access Control Model to Support Collaborative Processes

Xuan Hung Le and Dongwen Wang, University of Rochester Medical Center

EnhancedRBAC is an implementation of a role-based access control system (RBAC) that expands upon prior work by explicitly including the concepts of workflow and collaboration. The system architecture is divided into three layers: a policy encoding layer, a policy interpretation layer, and an application layer.

In particular, this paper evaluates the suitability of using EnhancedRBAC to organize training in association with New York State's HIV Clinical Education Initiative (CEI). The CEI requires coordination and cooperation between a number of different entities, which provide training materials for training on different topics, in different locations, and in person versus online; therefore employees must be able to access relevant training materials as a function of their role, the time, the training topic and location, etc. The authors acquired data on the current training system (CEI admin), acquired data for a reference standard by experts, and implemented EnhancedRBAC for CEI. The EnhancedRBAC performed against the current CEI system with a kappa of 0.8–0.89, and performed against the reference standard with a sensitivity of 97–100% and a specificity of 100%.

Tragedy of Anticommons in Digital Right Management of Medical Records

Quanyan Zhu, Carl Gunter, and Tamer Başar, University of Illinois at Urbana-Champaign

There are strong reasons for promoting EHRs, including the portability of data, increased timeliness, and reduced costs, but security and privacy concerns lead to the protection of EHRs via DRM. Different parties and institutions who contribute to the formation of different parts of the EHR have different technological, economic, and policy setups, leading to active or passive obstacles to sharing data (Tragedy of the Anticommons).

The authors apply cooperative and non-cooperative game theory to show that coordination is needed in order to achieve solutions optimal for multiple parties. They use the Shapley Value to model the value of sharing data and adding it to the EHR, which takes as input the cost of granting access, and the institution of a fee for accessing data.

On XACML's Adequacy to Specify and to Enforce HIPAA

Omar Chowdhury, The University of Texas at San Antonio; Haining Chen, Purdue University; Jianwei Niu, The University of Texas at San Antonio; Ninghui Li and Elisa Bertino, Purdue University

The requirements of HIPAA, which govern the collection, access, and storage of medical data for the privacy of the patient, are usually specified by researchers in a formal way. The actual implementation of HIPAA in organizations depends on internal policies and implementations. Meshing organizational formalisms and research formalisms usually requires manual inspection, but it would be ideal to automate this process.

The paper studies the applicability of OASIS's eXtensible Access Control Markup Language (XACML) to this problem. XACML is largely sufficient for HIPAA's needs, but it is primarily stateless, and therefore is unable to meet some of HIPAA's needs: obligations (e.g., must notify within 30 days), event history, policy-directed attribute retrieval, and policy-directed policy retrieval. Other HIPAA needs not met by XACML are the support for subjective beliefs (e.g., an employee believes that disclosure of PHI is necessary for law enforcement purposes) and reference to other policy rules or laws. The authors suggest high level extensions that could be made to XACML to support HIPAA—for example, they suggest using the obligation model of Li et al. 2010.

Access Control Hygiene and the Empathy Gap in Medical IT

Yifei Wang, Sean Smith, and Andrew Gettinger, Dartmouth College

Good access policies for EMRs are important in order to prevent the unauthorized access or alteration of PHI and facilitate access when necessary for the patient's health. The inaccurate or hurried creation or update of access control policies can lead to poor "access control hygiene" over time.

The paper presents a study exploring the "empathy gap" with respect to access control hygiene—the gap between what people think or decide when they reason in the abstract (about themselves or others) versus what people think or decide when they are directly embedded in a situation.

The study was run with 164 participants (78 experimental, 86 control), all medical employees. Overall, the study found that the experimental group (those using an EMR) made statistically more permissive decisions than those in the control group (creating access control policies), although there was noticeable variation in the results produced by the thirteen questions.

Audit

Summarized by Sang-Yoon Chang (email: chang6@illinois.edu)

Secure Logging and Auditing in Electronic Health Records Systems: What Can We Learn from the Payment Card Industry

Jason King and Laurie Williams, North Carolina State University

Jason King began the talk with the comparison between the health information technology (HIT) and the payment card industry (PCI). Both industries are similar in that they involve the management of sensitive, protected information, but they differ in the consequences for breaches; the breach of HIT data is more difficult to recover than PCI data.

In the approximately 75-page PCI data security standard requirements document, Jason King emphasized the complete nature of PCI requirements. As for health care industry enforcement, despite the Health Information Technology for Economic and Clinical Health (HITECH) Act of 2009, which provides incentives for health care professionals and hospitals to address and incorporate security and privacy for protected health information (PHI), the requirement was substantially less thorough than that of the PCI.

After discussing the necessity for stronger requirements and practice for PHI security (by presenting a scenario where a doctor conspires with an administrative user to grant an unauthorized privilege), Jason King concluded the talk with future work such as analyzing the effectiveness of securing PCI and vulnerabilities that are unique for HIT (and not applicable in the PCI industry).

One person mentioned that PCI took a lot of work and time to get to this state (in terms of security) and asked Jason King's opinion about whether he thinks the health care industry will go through similar experience for achieving security. Jason King responded that health care is about six years behind of the PCI industry. Somebody else asked how comparable the data storage and processing of the two fields are, to which Jason King replied that one distinction is that financial records by banks lack traceability. Another comment was that, even though many people who are involved in health care security are also involved in PCI, the two industries differed in the threat source; PCI encounters and targets hackers while health care seems to be less focused on that threat.

Enabling Robust Information Accountability in E-Healthcare Systems

Daisuke Mashima and Mustaque Ahamad, Georgia Institute of Technology

Daisuke Mashima discussed three trends for health care information: the transition from paper-based records to electronic records, the "meaningful use" incentive program (i.e., HITECH in 2009) to incorporate security, and initiatives for large-scale health information exchange. He moved

on to discussing threats such as the misuse of health data, data breach by insiders (specifically discussing a case in 2006 in Cleveland where a clinic employee sold information to a medical claims company, which then filed false claims to Medicare), and malware attacks.

To counter such threats, the authors propose incorporating accountability for electronic health record (EHR) sharing and usage, providing patients with verifiable evidence about who gets involved in the sharing path of EHR, and, hopefully, discouraging malicious and inappropriate handling by insiders. After presenting related work from other researchers, Daisuke Mashima presented their work that builds on previous work of a user-centric monitoring agent. He explained the architecture using a monitoring agent outside of the health care system and showed how the scheme works when updating ("accountable update") and using ("accountable use") the health information. The work also incorporates an accountability tag, metadata attached to the health data that will enable them to trace the sharing path. Lastly, Daisuke Mashima presented the prototype implementation, its integration in NHIN Direct, and the time overhead analysis (from which they concluded that the overhead for robust accountability is acceptable).

Someone raised the question about motivation to implement the scheme (and possibly to enforce it as a legal requirement). Daisuke Mashima responded that there may be a motivation for the organization to keep track of the flow of the data, since it owns the data.

Accountings of Relationships

Joseph Lorenzo Hall, New York University; Benedicte Callan, University of Texas at Austin; Helen Nissenbaum, New York University

Joseph Lorenzo Hall began his talk by discussing the HIPAA privacy rule based on fair information practices (FIP) and the emphasis on learning how the health information is used (in addition to what information is in the record). There are two mechanisms to track the use of health information: accounting of disclosure (AOD) and access report (AR), the latter proposed by Health and Human Services' Office for Civil Rights. From public comments that display frustrations/use cases, the authors discussed the problems associated with AOD and AR, including the fact that they are not widely used and the scaling problem. As an alternative, Hall presented accounting of relationships (AOR), which aggregates AODs and is essentially a data structure of aggregate information flow. He provided visualizations of the AOR concept using two types of diagram, the Sankey diagram (which shows network flows) and the Alluvial diagram (which shows flows over time). Hall concluded the talk with foreseeable challenges, including privacy in low-node flows and vocabulary compatibility across various institutions.

How were public comments sampled for analysis? They were randomly sampled. Someone commented that, from a sample access report that was presented during the talk, the records seem rhetorical, ineffective, and not universal. He also pointed out the challenge of deciding which systems are important and effective.