Report on the Sixth ACM Workshop on Privacy in the Electronic Society (WPES 2007)

Adam J. Lee
University of Illinois at Urbana-Champaign
adamlee@cs.uiuc.edu

Ting Yu
North Carolina State University
yu@csc.ncsu.edu

1 Workshop History and Overview

The world is transforming into an electronic society where almost every aspect of our lives is increasingly computerized and interconnected. Such a transformation has profoundly changed the scope, the scale and the level of automation for information collection, storage, analysis and dissemination. It, on the one hand, has and continues to enable new and better services. On the other hand, it inevitably increases the degree of privacy concerns.

The ACM workshop on Privacy in the Electronic Society (WPES) is dedicated to the discussion of problems related to privacy in today's global interconnected society. Since its establishment in 2002, WPES has been held in conjunction with the ACM Conference on Computer and Communications Security (CCS), and become an active forum for researchers and practitioners from both academia and industry to present novel research on the theoretical and practical aspects of electronic privacy, as well as experimental studies of fielded systems. Considering the broad implication of privacy, WPES welcomes submissions on a wide range of topics of interests, and encourages discussions and collaborations among researchers from multiple disciplines. As a consequence, each year the workshop attracts not only submissions with technical solutions from computer science's perspective, but also those from social science, laws and economics.

The Sixth WPES was held in Alexandria, Virginia on October 29, 2007. The workshop received 48 submissions, and accepted 9 full papers and 7 short papers. More than 40 participants joined the workshop.

2 Technical Program

The technical program for this year's workshop included three sessions for full-length research submissions, and a session devoted to short papers. Each of these papers is available for download from the ACM Digital Library.

2.1 Anonymous Communications

Anonymous communications concern about the design and analysis of protocols that protect the identities of the senders and receivers of Internet communications. Notable approaches include mix networks, DC-net, Freenet, Onion Routing and Crowds. Many anonymous systems have also been developed and deployed (e.g., Tor, Freenet, AP3 and GNUnet).

The first paper in this session was entitled "Probabilistic Analysis of Onion Routing in a Black-box Model." It quantitatively studied how, in an onion routing network, an attacker may gain more information of a user's identities when he possesses knowledge of the user's probabilistic behavior, especially when compared with that of other users. In particular, the paper observed that a user's anonymity is weakened either when the destinations that others visit are least likely visited by the user, or when others always visit the destination that the user chooses. The paper rigorously defined a probabilistic model to characterize the severity of privacy compromise due to the above observation. Though focused on onion routing, the paper's black-box model can be adapted to analyze other anonymous communication protocols.

The next paper entitled "Low-Resource Routing Attacks Against Tor" explored attacks against Tor, an anonymous communication system. Though most anonymous communication protocols can be shown in theory to provide strong protection of user privacy, when they are deployed, we often has to consider performance to make it more usable and practical. One such mechanism in Tor is preferential routing, where high-performance routers may be more likely to be selected. The paper showed an attack where attacker-controlled routers falsely claim to be of high performance so that they are chosen more often to appear in preferential routings. An attacker only needs to control only a few routers to significantly compromise the privacy of users. The paper suggested to use reputation mechanisms to verify the performance claims from

routers. This paper was a typical example of the intrinsic tradeoff between security and system utility. Similar observations are also found in data anonymization when anonymization algorithms are optimized to improve data utility.

The last paper in this session was entitled "Enhanced Privacy ID: A Direct Anonymous Attestation Scheme with Enhanced Revocation Capabilities." Direct Anonymous Attestation (DAA) is a technique for the remote authentication of a Trusted Platform Module (TPM) while preserving the user's privacy. Previous work on DAA requires the private key of a compromised TPM to be revealed before it can be revoked. This paper offered an improvement which revokes a compromised TPM without the need to know its private key. This scheme presented not only offers stronger privacy protection but also provides the same security guarantee as existing work on DAA.

2.2 Privacy in Distributed Systems

The first paper presented in this session was entitled "Making P2P Accountable without Losing Privacy" and presented an interesting solution to the problem of so-called *free riders* in peer-to-peer systems. The authors described a novel use of anonymous e-cash in which nodes that provide services to the network (i.e., share files) are rewarded with fungible credits that can later be used to purchase services from arbitrary nodes in the system. Since data must be purchased, rather than donated by altruistic network participants, nodes are required to share *at least* as much data as they download. The protocols presented in this paper enable users to be held accountable for their actions without compromising their privacy by explicitly linking them to their downloads.

The paper entitled "Improved User Authentication in Off-The-Record Messaging" addressed how usability concerns can undermine the security and privacy properties provided by OTR. Specifically, users that are unfamiliar with the basics of public key cryptography can make unsafe choices during connection establishment that allow a man-in-the-middle to observe and modify their supposedly private conversations. The authors then design and implement a modification to the OTR authentication and key exchange protocol that allows two users to safely and correctly establish an OTR session simply by knowing the same shared secret. This secret can be established either during an out of band protocol (e.g., meeting at a conference) or by providing hints to one another over an insecure channel (e.g., "What movie did we watch last week?"). An evaluation of their protocol shows that it is relatively inexpensive and correctly prevents MITM attacks.

The last paper presented in this section was titled "Single-bit Re-encryption with Applications to Distributed Proof Systems." The authors show that the use of *any* traditional public-key cryptography to protect information flowing through a distributed proof system introduces a covert channel that can be used to compromise the confidentiality of private facts. They then propose a single-bit re-encryption primitive, based on the Goldwasser-Micali cryptosystem, that eliminates this problem. Discussions at the workshop revealed that, in certain circumstances, malicious parties can collude out-of-band to infer the truth values of confidential facts using an attack similar to that which motivated this work. However, since these types of distributed proof systems are largely designed for use in pervasive computing spaces, their only means of communication is likely to be the proof protocol itself. As a result, eliminating the covert channel discovered in this paper is sufficient to prevent the inference of confidential facts.

2.3 Short Papers

Three of the papers presented during the short papers session were focused on the issues surrounding advanced authorization and authentication systems. "Enhancing Privacy in Identification Management Systems" addressed the issue of increasing user privacy in systems such as Microsoft's Cardspace. The authors showed two ways in which these types of systems can be extended to support privacy-enhanced claims (e.g., $age > 18$), rather than requiring the disclosure of raw claim data (e.g., disclosing the exact value of the *age* claim). The paper entitled "Harvesting Credentials in Trust Negotiation as an Honest-But-Curious Adversary" showed that a malicious party in a trust negotiation protocol can strategically alter the path of any negotiation to learn *each* of the victim's credentials that he is authorized to see. This attack requires no deviation from the underlying trust negotiation protocol. Lastly, the paper entitled "Information Carrying Identity Proof Trees" shows how previously-computed proof trees can be safely reused in advanced policy frameworks. This work provides a foundation for making these types of systems more efficient, as the cost of generating a particular sub-proof can be amortized over multiple uses.

The remaining papers from this session detail various ways in which information disclosure can reduce or increase user privacy in online systems. "Self-monitoring of Web-based Information Disclosure" presents a study of several visualization techniques that allow users to examine how their Internet search patterns vary over time.

The authors hypothesize that such visualizations may help users self-regulate their Internet usage and avoid disclosing too much personal data to online services. "Distance-Preserving Pseudonymization for Timestamps and Spacial Data" provides a technique through which data points in one- and two-dimensional spaces can be anonymized while still preserving the ability of third parties to compute the distances between points. Such techniques could be helpful in situations which mutually-distrustful partners wish to collaboratively monitor intrusion detection system logs. The paper entitled "Does Additional Information Always Reduce Anonymity?" shows that learning additional information regarding user input patterns or observed communications does not always help an attacker link the inputs and outputs of a threshold pool mix. This interesting result stems from a widespread belief that an attacker's uncertainty is the same as Shannon's notion of conditional entropy. The authors present a counterexample proving their claim and clarify the differences between these two concepts. The final paper presented in this session was entitled "Disappearing For A While—Using *White Lies* in Pervasive Computing." This paper addresses the socio-technological issues surrounding the use of lies in location detection systems. The authors argue that users will want to "disappear" once in a while to preserve their privacy and present a framework through which users can lie about their present location without corrupting the context of the pervasive computing system.

2.4 Privacy preservation and Social Issues

The paper entitled "Private Web Search" presented a discussion of the ways in which simply searching the web can lead to violations of users' privacy. This paper is well-motivated in light of the August 2006 release of "anonymized" AOL search engine logs that, in many cases, could still be used to identify the users who made certain groups of queries. It is argued that information that can be used to potentially identify a user or correlate their searches is leaked at the network level via IP addresses and the like, in HTTP headers via cookies and software version information, at the HTML level through JavaScripts or timing attacks, through the search terms used by an individual (e.g., names or addresses), and by active components in a web page. The authors then describe PWS, a tool that they have developed to help protect user privacy while searching the web.

In the paper "Towards Understanding User Perceptions of Authentication Technologies," the authors report on the results of a preliminary survey conducted to assess users' beliefs about various authenticators. In particular, the authors sought to evaluate users' familiarity with and preferences for various authentication technologies, as well as the users' perceptions of the usefulness, acceptability, security, and privacy of these technologies. This study found that users overwhelmingly preferred technologies that they were familiar with (e.g., passwords and fingerprint scans) over technologies that they did not understand (e.g., digital certificates). However, the results also indicated that users were willing to adopt new technologies, but had concerns about conflicts with religious beliefs and the misuse of data collected during the use of these technologies (e.g., the use of RFID for location tracking). Understanding these types user beliefs can help system builders better evaluate the impact of new technologies.

The last paper presented in this workshop was entitled "PriPAYD: Privacy Friendly Pay-As-You-Drive Insurance." Pay-As-You-Drive (PAYD) insurance policies compute an individual's insurance premium based upon factors such as the distance driven, the types of roads used, the times at which trips occurred, and whether speed limits were obeyed by the driver. Although PAYD policies are becoming more popular around the globe, these systems often use GPS technologies to record the fine-grained information necessary to compute these personalized bills. In this paper, the authors propose a privacy friendly alternative to existing PAYD models that allows the same type of fine-grained billing to take place without disclosing exact location information to any third parties. In their solution, a trusted black box under the control of the insurance company collects GPS data locally and sends only aggregate statistics to the insurance company for billing purposes. Policyholders can verify that only aggregate data is transmitted to the insurance company, while the insurance company can verify that policyholders do not tamper with the data transmitted for billing. This system shows that with careful design, the benefits of PAYD insurance can be realized without the loss of privacy.

Acknowledgments

The success of WPES 2007 relied on the volunteer efforts from all the members of organizing committee, the program committee, and the external reviewers. We have four senior members serving on the steering committee who initiate the workshop every year and provide advice to the organization of WPES. They are Pierangela Samarati, University of Milan, Italy, Sabrina De Capitani di Vimercati, University of Milan, Italy, Sushil Jajodia, George Mason University, USA, and Paul Syverson, Naval Research Laboratory, USA.