

Research Statement

William H. Winsborough

My current research interests are in computer security and privacy in distributed systems and networks, with an emphasis on policy-based techniques. I am particularly interested in techniques for managing the sharing of resources across multiple organizations or the whole Internet while protecting them from misuse. In such decentralized environments, the diverse interests of all participants must be factored into system behavior. These interests often require properties such as confidentiality, integrity, and availability of resources, including the resources of the security system itself, such as policies and credentials.

Recent Accomplishments

One of the chief requirements of any security system is preventing unauthorized access to resources. This task, access control, is particularly challenging in large systems in which authority emanates from many sources, such as individuals, organizations, or other resource providers. Traditional authorization mechanisms used in databases and operating systems are based on authenticating the identity of resource users. This approach works well when a prior relationship has been established between the user and the resource provider. However, when authorization systems span multiple organizations or grow spontaneously, as in the Internet, this approach breaks down.

For the last several years, I have focused on developing access-control techniques that enable users and resource providers to establish appropriate trust relationships automatically, without prior contact. My basic approach is to establish trust by making use of attributes of the individuals on whose behalf software entities operate—attributes such as the individuals' roles within their home organizations, relationships between organizations, individual or organizational qualifications, or other assessments by third parties of their trustworthiness. In the trust management approach, policy statements about such attributes can be carried in cryptographically signed credentials that resemble letters of recommendation through which issuers attest to precise and limited trust relationships.

Trust Management Working with colleagues at Stanford, I designed a family of efficient, expressive policy languages called Role-based Trust-management (*RT*) [4]. *RT* supports a rich model of delegation, and provides a natural, highly scalable administrative model. In *RT*, a credential is just a policy statement cryptographically signed by its issuer. One of the issues inherent in such a system is the availability of credentials to the authorization-decision engine. In the systems of interest, credentials are issued in a decentralized, non-hierarchical manner and, particularly in open systems, credential storage is almost inevitably distributed. Reflecting this, credentials in *RT* systems are stored with either their issuer or their subject. I developed techniques to discover and evaluate relevant credentials when an authorization query is posed [5, 7]. I developed a type system for credential storage that ensures relevant credentials can be found when needed, while allowing flexibility about where in the distributed system credentials are stored. This enabled me to prove completeness of the authorization-decision engine in the context of distributed credential storage, which complements my somewhat less novel but equally important soundness results. The evaluation algorithm collects only credentials that are relevant to the authorization decision, and ensures that queries are answered efficiently.

A further, related problem I solved addressed the fact that policies change. I designed mechanisms for scalable, rapid, and highly available publication of credential status (revocation/revalidation) that

are suited to administrative models in which there are many credential issuers [1]. By using our mechanisms, neither credential issuers nor parties relying on credentials need to trust publishers of credential status data, but rather can verify for themselves that the publisher is correctly posting up-to-date status data.

Automated Trust Negotiation Credentials may be sensitive, for instance, if they contain financial or medical data. The subject of such a credential has an interest in protecting it, yet needs to present it to gain authorization for certain resources. An inherent problem here is to enforce the confidentiality requirement, while disclosing the credential to appropriate entities when needed. In some of my earliest work in computer security, I introduced a simple approach to this problem called automated trust negotiation (ATN), in which participants establish trust in one another through cautious, iterative, bilateral disclosure of credentials [12]. The distinguishing characteristic of ATN differentiating it from most other trust-establishment schemes is that credentials themselves are treated as protected resources. Several separate research groups have now contributed to the rapidly growing literature on ATN.

I developed an ATN protocol, called TTG for the trust-target graph that represents protocol state [10, 9]. TTG supports an *RT* credential language that has delegation and distributed storage of credentials. In addition to controlling the transmission of credentials, TTG was designed to protect the information that one has sensitive attributes. I proposed several candidate, information flow-theoretic requirements for the safety of an ATN protocol, analyzed their relationships, and showed why one of them should be preferred [11]. Many protocols in the literature clearly fail to meet that requirement; however, I showed that TTG satisfies it [8].

Policy Analysis A central feature of authorization systems that allow scalable sharing of resources across multiple organizations is delegation. For instance, organization A may delegate to organization B the authority to determine which employees of organization B are assigned to a joint project. Thus, policy authors delegate authority to one another for defining portions of the policy. In this context, policy authors require assistance understanding the effect of their policy statements. I designed policy analysis techniques that answer important questions about such things as who can potentially gain access that perhaps should not (safety) and who may not be guaranteed access that they require for correct operations (availability). The analysis answers questions about the implications of an authorization policy, both now and after other policy authors have exercised their authority to change the policy in arbitrary ways. I showed that, unlike in many access control systems, the security analysis problem in *RT* is decidable. I also showed that many security analysis instances can be solved efficiently and gave complexity bounds in the intractable cases. This work has recently been accepted to appear in the *Journal of the ACM* [6, 3].

The analysis provides a means for ensuring security requirements are always met by the authorization system as the policy evolves. One of the design principles of trust management is that a policy statement cannot be removed or modified, except by its author. While assuming that a specified set of trusted individuals will not change their statements, the analysis determines whether a security requirement could nevertheless be violated in the future. By using the analysis to screen candidate changes before making them, the trusted individuals can ensure security requirements are never violated even while sharing control over the policy as a whole with untrusted or semi-trusted entities.

Recent and Current Funding

All of the research described above had federal funding. During my three years at NAI Labs, I was the PI on DARPA projects valued at about \$500K per year. I wrote the proposals, managed the projects, and was the technical lead. I am currently a co-PI on a five-year, medium-sized NSF ITR project awarded in 2003 to advance ATN research; project co-PIs are at BYU, Purdue, Stanford, UIUC, and USC.

Earlier Research Contributions

I contributed to several areas in programming languages, including program analysis, language implementation, and programming-support tools. I was among the first investigators to apply abstract interpretation to the static analysis of logic programming languages, a field which has received considerable attention. I developed mathematically sound dataflow analyses and compiler optimizations that improve the time and space performance of logic programs and concurrent logic programs. I also developed analyses of synchronization properties of concurrent programs, such as analyses that verify that a program is deadlock-free, and tools that support automatic granularity control and compile-time scheduling. Much of this work was funded by the NSF.

Future Research Directions

There are several lines of research I plan to pursue over the next few years. Below I sketch three of these. The thread that runs through these and other topics I hope to work on is the combination and enforcement of policy statements that express the interests of multiple authors with respect to confidentiality, integrity, and resource usage.

Security Analysis, Insider Threat Assessment, and RBAC The work I have already done lays a solid foundation for security analysis and opens up several new research possibilities whose culmination will be practical policy authoring tools. Many of the analysis questions in the class that is generally intractable would be extremely beneficial to answer. Examples include the question whether it is always true that only managers have access to a particularly sensitive document or that all employees have access to a company's internal web pages. Recall that security analysis asks whether such properties are invariant across policy-state changes made by untrusted parties. A related question is whether invariants are preserved across changes made by collections of possibly colluding trusted parties. I call this problem *insider threat assessment*, and am currently working to solve it. In addition, I am studying the application of these analyses to the popular authorization model, Role-Based Access Control (RBAC).

Policy-driven Email Services Another line of research involves applying attribute-based trust management techniques to managing email and controlling spam [2]. The idea is to provide each party participating in email transfer the ability to control their participation by writing policy based on a wide range of characteristics of the transaction and its participants. For instance, a message might be rejected if the sender is on a blacklist of known spammers, however it still might be accepted if a monetary bond were posted which the recipient could, at his discretion, seize upon receiving the message. The big goal is to provide a framework in which a wide range of techniques can be applied and controlled. Several interesting technical issues arise. For instance, the contents of a blacklist may be private. However, the policy mentioned above gives the sender a way to check whether he is on the blacklist by the correct functioning of the system itself: if a bond is seized, he must be on the blacklist. I plan to provide the author of such a policy practical means to ensure confidential information such as this is not leaked.

Information-Flow Control While in common use, access control policies are deficient in that they do not govern the further dissemination of information after it has been released to an authorized reader. Essentially, the reader must be trusted not to disclose the information to an unauthorized third party. A more comprehensive approach is instead to employ *information-flow* policies. These are policies that dictate not only whether a principal may read a document, but also whether it may be shared with someone else. Significant work has been done over the past several years developing program analysis techniques that determine whether a program satisfies information-flow policies. However, that work does not accommodate policies that evolve. In particular, revocation of rights is not supported. By working with colleagues who study on-the-fly software upgrades, I plan to devise techniques to update security policies safely, while they are in use, to prevent unintended leakage of confidential information.

References

- [1] M. T. Goodrich, M. Shin, R. Tamassia, and W. H. Winsborough. Authenticated dictionaries for fresh attribute credentials. In *First International Conference on Trust Management*, pages 332–347. Springer-Verlag, May 2003.
- [2] Saket Kaushick, Paul Ammann, Duminda Wijesekera, William H. Winsborough, and Ronald Ritchey. A policy driven approach to email services. In *IEEE 5th International Workshop on Policies for Distributed Systems and Networks*, pages 169–178. IEEE Press, May 2004.
- [3] Ninghui Li, John C. Mitchell, and William H. Winsborough. Beyond proof-of-compliance: Security analysis in trust management. To appear in *Journal of the ACM*.
- [4] Ninghui Li, John C. Mitchell, and William H. Winsborough. Design of a role-based trust management framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 114–130. IEEE Computer Society Press, May 2002.
- [5] Ninghui Li, William H. Winsborough, and John C. Mitchell. Distributed credential chain discovery in trust management (extended abstract). In *Proceedings of the Eighth ACM Conference on Computer and Communications Security (CCS-8)*, pages 156–165. ACM Press, November 2001.
- [6] Ninghui Li, William H. Winsborough, and John C. Mitchell. Beyond proof-of-compliance: Safety and availability analysis in trust management. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 123–139. IEEE Computer Society Press, May 2003.
- [7] Ninghui Li, William H. Winsborough, and John C. Mitchell. Distributed credential chain discovery in trust management. *Journal of Computer Security*, 11(1):35–86, February 2003.
- [8] William H. Winsborough and Ninghui Li. Safety in automated trust negotiation. Submitted to *ACM Transactions on Information System Security*, May, 2004.
- [9] William H. Winsborough and Ninghui Li. Protecting sensitive attributes in automated trust negotiation. In *Proceedings of the ACM Workshop on Privacy in the Electronic Society*, pages 41–51. ACM Press, November 2002.
- [10] William H. Winsborough and Ninghui Li. Towards practical automated trust negotiation. In *Proceedings of the Third International Workshop on Policies for Distributed Systems and Networks (Policy 2002)*, pages 92–103. IEEE Computer Society Press, June 2002.
- [11] William H. Winsborough and Ninghui Li. Safety in automated trust negotiation. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 147–160. IEEE Computer Society Press, May 2004.
- [12] William H. Winsborough, Kent E. Seamons, and Vicki E. Jones. Automated trust negotiation. In *DARPA Information Survivability Conference and Exposition*, volume I, pages 88–102. IEEE Press, January 2000.