Guest Editorial: Special Issue on Computer and Communications Security

This special issue contains extended versions of articles selected from the program of the *15th ACM Conference on Computer and Communications Security (CCS'08)*, which took place October 27 to 31, 2008 in Alexandria, Virginia (USA). This annual conference is a leading international forum for information-security researchers, practitioners, developers, and users to explore cutting-edge ideas and results, and to exchange techniques, tools, and experiences. Its mission is to promote and share novel research from academia, government, and industry covering all theoretical and practical aspects of computer security as well as case studies and implementation experiences. The selected articles represent the broad scope of the conference. Two articles focus primarily on what we can learn from novel and existing attacks on security, and the other two articles focus on the development of improved mechanisms to provide and assure security. The topics range from programming exploits to large network designs, from cryptographic advances to formal techniques to verify the use of cryptography in practical protocols.

The four articles in this special issue were invited for submission from the fifty-one papers presented at CCS'08. These were selected from 280 papers submitted to the conference. The submissions to the special issue were required to contain at least 25% new material to differentiate the journal articles from the conference papers. All the journal submissions went through an additional thorough review process (the same review process as any submission to this journal) to further ensure their quality.

The first article, "Return-Oriented Programming: Systems, Languages, and Applications" by Ryan Roemer, Erik Buchanan, Hovav Shacham, and Stefan Savage, describes a technique to attack programs without injecting code by chaining instruction sequences already in a program's address space. The technique they demonstrate challenges the assumption that preventing the introduction of malicious code can prevent malicious computation.

The next article shows an advance on the defense side of program security. In "Verified Cryptographic Implementations for TLS", Karthikeyan Bhargavan, Ricardo Corin, Cédric Fournet, and Eugen Zălinescu show how to develop a small functional implementation of TLS and then use novel and existing model-extraction and verification tools to provide security guarantees at both the symbolic and computational levels. The end result is the first automated verification of executable code using standard cryptographic assumptions.

Next, Jan Camenisch and Thomas Gross, in their article "Efficient Attributes for Anonymous Credentials," address a different aspect of cryptographic protocols, how to practically provide just the authorization needed. A novel approach permits much more efficient proofs of possession of even a large combination of attribute credentials. This makes it much more practical to anonymously show that one has needed attributes, for example, based on the set of credentials represented on a digital identity card or other computationally limited device.

Finally Prateek Mittal and Nikita Borisov also examine anonymity, but for large networks rather than small devices. "Information Leaks in Structured Peer-to-Peer Anonymous Communication Systems" analyzes the node lookup mechanisms in published systems. They show that robustness to active attacks can only be improved with an increased vulnerability to passive attacks, thus limiting the effectiveness of these approaches to scaling anonymous communication using peer-to-peer networks.

We are thankful to all who contributed to this special issue: the referees for their helpful and insightful comments, the authors for their hard work and their patience, and Editor-in-Chief Gene Tsudik for his suggestions and support.

—Paul Syverson
Somesh Jha
Guest Editors