Guest Editors' Preface

This issue of the Journal of Computer Security is drawn from papers presented at the 1991 IEEE Symposium on Research in Security and Privacy, Oakland, California, USA, May 1991. These annual symposia have been sponsored since 1980 by the IEEE Computer Society Technical Committee on Security and Privacy. The Symposium is a forum for reporting research results in computer security from academic, industrial, and government laboratories.

Several papers presented at the 1991 Symposium were invited to be submitted to a special issue of the Journal. Submitted papers were revised for Journal publication and subjected to the normal rigorous review process of the Journal. This issue contains five papers selected for publication through this process.

In "An Analysis of Covert Timing Channels" John C. Wray notes that some covert channels are not easily characterized as storage channels or timing channels, but have aspects of both. Wray presents a method for constructing channels with a timing nature from a knowledge of a system's asynchronous behavior. Noting that processes can use event streams as clocks, Wray finds that a timing channel may be exploited when the receiver has access to two event streams, one of which is modulated by the sender and one of which is used as a reference clock. His technique enumerates all "clocks" in the system and then examines these pairwise to determine whether each pair of clocks can be exploited as a channel. This approach was used in the covert channel analysis for the VAX VMM, a virtual machine monitor designed to meet the requirements for Class A1 of the DoD Trusted Computer System Evaluation Criteria.

Wei-Ming Hu's paper, "Reducing Timing Channels with Fuzzy Time," describes techniques that can be applied to reduce the bandwidth of covert timing channels such as those identified in Wray's paper. The techniques were developed to address high-speed hardware timing channels, such as bus contention channels. Hu demonstrates that conventional techniques, such as closing the channel directly, auditing the use of the channel, and making the channel noisy, either do not work or are highly impractical for these types of channels. To make use of a covert channel, a process can use a stream of I/O operations to simulate an interval timer, or it can reference the system clock. Hu's technique defeats such clocks by reducing the accuracy and precision of the system clocks and by randomly altering the timings of I/O operations. Hu's techniques have been implemented in the VAX VMM security kernel.

The paper by James W. Gray III, "Toward a Mathematical Foundation of Information Flow Security," describes a probabilistic trace-based framework that can be used to model nondeterministic computer systems. Gray uses this framework to define PNI, a new probabilistic model of information flow, and AFM, a formalization of McLean's probabilistic model, FM. He goes on to show that PNI is weaker than AFM, but still strong enough to prevent information flow from high-level users to low-level users. He discusses the reasons for the difference between PNI and AFM, gives verification conditions for both models, and compares PNI to other security models. Probabilistic security models, such as the ones Gray discusses in this paper, are needed because nondeterministic systems may exhibit probabilistic covert channels that are not ruled out by standard computer security models.

In "SPX: Global Authentication Using Public Key Certificates" Joseph J. Tardo and Kannan Alagappan describes SPX, a reference implementation of a distributed authentication service architecture intended for open network environments. SPX supports mutual authentication of applications in arbitrarily large networks with multiple, mutually suspicious jurisdictional authorities. In contrast to Kerberos, which uses only secret keys, SPX uses a combination of public key and secret key techniques. The use of public key techniques eliminates the need for on-line trusted key distribution centers and allows SPX to scale well since there's no need for on-line, globally trusted authorities. Additional applications and capabilities are planned, e.g., to allow the selective control of access to remote files.

Paul Syverson's paper, "The Use of Logic in the Analysis of Cryptographic Protocols," surveys logics for cryptographic protocol analysis and gives an insightful analysis into possible sources of confusion and contention that have surrounded this topic in the past. He argues there is no significant syntactic distinction between those cryptographic protocol logics that are based on knowledge and those that are based on belief, and he argues that all cryptographic protocol logics would benefit from an independently motivated formal semantics. As an example of his latter point, he shows how Abadi and Tuttle's semantics for the logic of Burrows, Abadi, and Needham resolves a debate over an alleged flaw in this logic put forward by Nessett. We expect that this paper will serve as a stimulus for further discussion rather than as a final word on the subject.

In conclusion, we thank the authors and reviewers of the papers invited for the Journal for their hard work. We thank the editors of the Journal for encouraging us to compile this special issue of the Journal and to present the best work reported at the 1991 Symposium to the larger readership of the Journal.

Teresa F. Lunt
John McLean