Editor's Preface

This double issue focusses on database security. Two strive for concurrency control without covert channels, one addresses the inference problem, and two propose more general access control policies. All of them are motivated by real-world concerns for efficiency and usability.

Concurrency control in a multilevel database involves decomposing multilevel transactions into single-level transactions, in such a way that correctness (serializability) is preserved, and there are no signalling channels due to higher-level activity or aborts affecting lower-level results. Data replication is necessary to prevent possible starvation of higher-level transactions. "Concurrency control in a secure multilevel database via a two-snapshot algorithm", by P. Ammann et al., calls for only two replicates per level, representing older consistent states of the database which are read by higher-level transactions. "Transaction management for multilevel secure replicated databases", by I. Kang, T. Keefe, assumes that each level of data is replicated for all higher levels, but transactions need only read local copies at their own level. The latter paper gives two transaction scheduling algorithms, one which is valid only for some partial orderings of sensitivity levels, and another, timestamp-based, for any partial ordering. The tradeoffs in this area are complex, and these papers have a deeper discussion of the issues.

Covert channels in multilevel relational databases can also arise from inferences due to functional dependencies. There are hidden, indirect associations among attributes that may violate sensitivity level assignment policies. "A fast algorithm for detecting second paths in database inference analysis", by T. Hinke et al., adapts a classical algorithm for join analysis to obtain an order of magnitude speedup over path-finding approaches used in other work. It is always a pleasure to see work that makes good use of related past advances.

Security policies for databases go far beyond label-based confidentiality. It is important to have conceptual apparatus for expressing policies that reflect the organizational constraints and flexibility characteristic of commercial enterprises. We have two papers that propose rich models for access authorization and control. "An extended authorization model for object databases", by E. Bertino et al., presents an authorization structure incorporating object inheritance, versions, and composite objects, as well as a user group hierarchy. It has a unique approach to implied and negative authorizations. "Merging models: integrity, dynamic separation of duty, and trusted data management", by L. Notargiacomo et al., focusses on the dynamic aspects of policies, and introduces a concept of controlled activity in the form of predefined transactions. It supports an interpretation of the Clark–Wilson model and the Brewer–Nash "Chinese Wall" policy for relational databases. Both of these papers pay careful attention to how their approaches can be implemented with straightforward modifications of existing systems.

Jonathan K. Millen