FORMAL METHODS IN CIS: HETEROGENEOUS DATABASES GUEST EDITORS' INTRODUCTION

MICHAEL N. HUHNS

Department of Electrical and Computer Engineering, University of South Carolina
Columbia, SC 29208 USA

MUNINDAR P. SINGH

Department of Computer Science, North Carolina State University
Raleigh, NC 27695-8206, USA

There has been much progress recently in cooperative information systems: new problems are being attacked, new techniques are being developed. As the field matures and as its applications become deployed in a wider range of domains and locales, there is increasing interest in formal methods to understand, analyze, design, verify, and validate cooperative information systems. Since the science of cooperative information systems (CIS) unabashedly borrows from several traditional disciplines, the techniques and formal methods that it applies are also of a wide variety.

This issue presents some of the best papers in this subfield of CIS. These papers are unified by their relationship to heterogeneous databases and information environments. Even within this subset of CIS, it is interesting to see methods whose traditional counterparts are in artificial intelligence, software engineering, and database semantics and query processing.

Lee, Madnick, and Siegel address one of the fundamental problems in CIS, namely, that of the semantic interoperability of information systems that arise in different contexts. Context is the holy grail of a number of sciences that deal with meaning. This paper proposes an approach that has many useful properties for enabling the exchange of information among diverse data sources. It models context and interoperation at the knowledge level, term coined by Allen Newell in artificial intelligence; it uses techniques pioneered by Ray Reiter, also in artificial intelligence. This enables an abstract characterization of the meaning of information independent of its specific realization in a particular system.

Lee and Wu also concern themselves with the interoperability of autonomous databases. They propose hyperrelations—extended from the relations of traditional database theory—as a uniform and succinct means to represent the local schemas

of the databases being interoperated. This facilitates detecting the several possible kinds of conflicts among local schemas. Further, they develop a hyperrelational algebra, which enables query transformation and optimization.

Florescu, Raschid, and Valduriez address the problem of query reformulation in a CIS. Query reformulation involves transforming an input query stated with respect to one interface into a set of queries that can be executed on the available information systems. Thus it involves both decomposition and result merging. This obviously presupposes semantic knowledge, which the present paper elegantly captures in the form of rewrite rules. The proposed approach also accommodates certain kinds of query optimization.

Even, Faase, and de By present an approach for formally specifying CIS using techniques that were first developed in the traditional settings of software engineering and object-oriented databases. They study the interaction of two orthogonal, well-known languages: TM (based on type theory and designed for object oriented database schemas) and LOTOS (based on process algebra and for distributed systems). This paper represents the kinds of heavy-duty formal specification techniques that we will ultimately need for CIS, and the kinds of extensions to traditional techniques those will have to embody.

We expect the papers of this issue to have a great impact on the formal methods in CIS. However, we fervently hope that these papers will not be the last word on their subjects, but will provide a solid foundation for the problems that remain unsolved or even unformulated today.