



Guest Editorial

Together with wireless communication technologies, portable computers and personal digital assistants (PDAs) provide a pervasive base for mobile computing. As the mobile computing technology matures, millions of people are becoming mobile users communicating with one another and accessing various information resources. Mobile computing involves mobility of users, hardware, software and data.

Besides advances in communications and hardware, achieving pervasive mobile computing requires innovative theories, paradigms and applications in data management. MobiDE (ACM International Workshop on Data Engineering for Mobile and Wireless Access) is a biannual workshop that aims at acting as a bridge between the data management and the mobile computing communities. This special issue includes extended versions of papers selected from the second in this series of workshops (MobiDE01) that took place in May 20, 2001, in Santa Barbara, in conjunction with SIGMOD.

The four papers included in this issue cover some of the most important aspects of the mobile data management field and can serve as a reference point for this exciting area. The issues tackled include resource discovery, mobile publish/subscribe, location-based services and personalized data dissemination.

In more details, the first paper by A. Friday et al. covers the issue of service discovery and querying in ubiquitous computing. The authors argue that existing service discovery protocols fail to address the requirements of ubiquitous applications that involve the simultaneous invocation and cooperation of multiple specialized location and interaction protocols. To address these limitations, a simple extensible meta-service discovery architecture that uses database techniques to unify service discovery protocols is presented along with its prototype implementation.

The focus of the second paper by Y. Huang and H. Garcia-Molina is on publish/subscribe systems which are systems that dynamically route and deliver events to interested users that have previously subscribed to the system. The paper discusses how such a publish/subscribe system can be adapted to operate in a mobile environment, where events can be generated by moving users or sensors and subscribers may request delivery at handheld and potentially mobile devices. In particular, the authors show how the system itself can be distributed across multiple, possibly mobile computers to achieve load balancing and how the system can exploit replication to handle message losses, failures and network disconnections.

The third paper by B. Zheng, W.-C. Lee and D.L. Lee deals with location-based services and in particular with mobile nearest-neighbor search in which a mobile user issues a query for retrieving all stationary objects located closest to it. In order to support such queries, an indexing scheme based on Voronoi diagrams is introduced. In addition, semantic cache is used for improving performance. Query scheduling policies are presented that address issues that arise in multi-cell environments where a client roams freely across different cells. Simulation results are presented to validate the proposed approaches.

Finally, the fourth paper by B. Ozen et al. addresses the important issue of personalization. The authors present an architecture for personalized information delivery from XML resources to mobile clients based on user profiles. For scalability, indexing of the user profiles instead of the documents is supported. In this way, all queries that apply to a specific document are executed in parallel by using a finite state machine (FSM) approach for parsing the document. Moreover, queries that have the same FSM representation are grouped together. The system includes a user-friendly graphical interface.

In closing, we would like to thank all those that contributed to this special issue: the editor-in-chief, Prof. Imrich Chlamtac, for his cooperation, the reviewers for their thorough comments that help in enhancing the quality of the papers and the authors for providing extended and improved versions of their workshop papers.

Panos K. Chrysanthis and Evangelia Pitoura
Guest Editors



Panos K. Chrysanthis is a Professor of Computer Science at the University of Pittsburgh and an Adjunct Professor at Carnegie Mellon University. He received his B.S. from the University of Athens, Greece (1982) and his M.S. (1986) and Ph.D. (1991) from the University of Massachusetts at Amherst. His current research focus is on mobile and pervasive data management. In 1995, he was a recipient of the NSF CAREER Award for his investigation on mobile data management. Besides journal and conference articles, his publications include a book and book chapters on transaction processing and on consistency in distributed, mobile databases and multidatabases. He is currently an editor of the VLDB Journal, and has served as general chair and program co-chair of several international workshops and conferences related to mobile computing.



Evaggelia Pitoura received a B.Sc. degree in computer engineering (1990) from the Polytechnic School of the University of Patras, Greece, and a M.Sc. (1993) and a Ph.D. (1995) degrees in computer science from Purdue University, USA. She is currently an assistant professor at the University of Ioannina, Greece. She has served on the Program Committee of more than 30 international workshops and Conferences and co-chaired three international workshops. For her work on mobile agents, she has received the best paper award at ICDE 1999. Besides journal and conference articles, her publications include a book on mobile computing. Her research interests include mobile and pervasive data management and more recently peer-to-peer computing.