

Report on MobiDE 2003: The 3rd International ACM Workshop on Data Engineering for Wireless and Mobile Access

Sujata Banerjee
Hewlett Packard Labs
Palo Alto, CA, USA
sujata@hpl.hp.com

Mitch Cherniack
Dept of Computer Science
Brandeis University
Waltham, MA, USA
mfc@cs.brandeis.edu

Panos K. Chrysanthis
Dept of Computer Science
University of Pittsburgh
Pittsburgh, PA, USA
panos@cs.pitt.edu

Vijay Kumar
Dept of CS and EE
University of Missouri, Kansas City
Kansas City, MO, USA
kumarv@umkc.edu

Alexandros Labrinidis
Dept of Computer Science
University of Pittsburgh
Pittsburgh, PA, USA
labrinid@cs.pitt.edu

The 3rd International ACM Workshop on Data Engineering for Wireless and Mobile Access (MobiDE 2003 for short) took place on September 19, 2003 at the Westin Horton Plaza Hotel in San Diego, California in conjunction with MobiCom 2003. The MobiDE workshops serve as a bridge between the data management and network research communities, and have a tradition of presenting innovations on mobile as well as wireless data engineering issues (such as those found in sensor networks). This workshop was the third in the MobiDE series, MobiDE 1999 having taken place in Seattle in conjunction with MobiCom 1999, and MobiDE 2001 having taken place in Santa Barbara in conjunction with SIGMOD 2001.

1 Workshop Overview

Our call for papers attracted 34 high-quality submissions from 7 countries, making the selection very competitive. Most papers were reviewed by 4 members of the Program Committee, with only a few papers receiving 3 reviews. After a discussion phase on a few papers with conflicting reviews, 9 full papers were accepted. This represents a highly competitive acceptance rate of 26.5%. In addition,

we accepted 5 short papers that represented promising works in progress.

One of the chosen papers, *Consistency Mechanisms for a Distributed Lookup Service supporting Mobile Applications* [7], by Christoph Lindemann and Oliver Waldhorst of the University of Dortmund was chosen as best paper. This paper, which presented a general-purpose distributed index lookup service for mobile devices, was subsequently invited to appear in a special issue of ACM SIGMOBILE Mobile Computing and Communications Review [3].

MobiDE 2003 also introduced a novel *student grant competition* open to all students with the purpose of subsidizing their attendance at the workshop, thanks to a generous donation from the National Science Foundation.¹ A total of 14 grants of \$600 and up were awarded to students in order for them to attend the workshop. Grants were given to 10 graduate students from the US, one undergraduate student from the US, and three graduate students from abroad. Including the student attendees, there were 44 registered participants for the workshop.

¹NSF Award #IIS-0334025 from the Information and Data Management Program provided the funds to support student travel to the workshop.

2 Workshop Schedule

The day-long workshop began early on September 19 with an informal breakfast, and was followed by opening remarks from Prof. Panos Chrysanthis of the University of Pittsburgh, the workshop's general chair. The keynote address: "*Data Management in Sensor Networks: Challenges and Opportunities*", was then presented by Dr. Wei Hong of Intel Research, Berkeley.

In his talk, Dr. Hong presented the TinyDB and ASK acquisitional query processing system out of Intel Labs and University of California, Berkeley. The talk presented the design and implementation of this in-network sensor database system and supporting toolkit, and then described and demonstrated a number of TinyDB applications, including an ecological application that used embedded sensor devices to track the mating and migration habits of birds. The keynote talk set up the stage for many of the following presentations in a variety of ways. It provided a different way of seeing and addressing issues common to mobile computing and sensor devices, such as power-awareness and power management, as well as problems related to data communication over wireless links. The remainder of the day was devoted to presentation of research papers, as described below.

Session 1: The first paper session was devoted to *Data Dissemination and Pervasive Computing*. This session included three long (30 minute) and one short (15 minute) talks. The first talk proposes context oriented programming (COP) which elevates *context* to a first-class construct [6]. The claim is that in ubiquitous computing environments, where products need to be adaptable and portable and yet still retain a simple code base, COP could provide some advantages. The second paper outlines the potential role that semantic techniques offer in solving some key challenges, including candidate service discovery, intelligent matching, service adaptation and service composition [10]. The short paper of the session described the DAYS architecture, which is designed to provide a flexible broadcast environment which allows clients to update the content of the broadcast [2]. This paper led to a long discussion about the

adoption of data dissemination techniques from the industry, given that there exists a significant body of work on this area (some of which was also presented in subsequent sessions at the workshop, prompting a continuation of the discussion). The final long paper of the session described a decentralized weighted voting scheme for managing replicated data in a mobile peer-to-peer system [12], which was one of the earliest papers in the area of mobile peer-to-peer networks. Weighted voting offers a familiar consistency model and supports on-line replica reconfiguration, which the authors argued makes it a good fit for applications in the pervasive computing domain.

Session 2: Location awareness in data, queries and users is one of the key characteristics of mobile computing and still an open and challenging research area. Thus the second paper session was devoted to *Location Awareness and Moving Objects* which included two long and three short talks. The first paper of this session proposed an efficient method to place geographical data items over broadcast channels that reduces access time for spatial range queries on them [17]. The second paper addressed the issue of answering spatio-temporal range queries when there is uncertainty associated with the model of the moving objects [14]. The authors proposed a framework based on the concept of trajectories to capture the spatio-temporal properties of moving objects and show how queries whose results are invalidated by changes in the database (environment) can be efficiently identified. The first short paper of the session proposed a low-cost, two-step location update/paging scheme in a macrocell/microcell network [15]. The savings in operating cost is obtained by conducting location updates only in the macrocell tier. The next paper proposed a data storage system for mobile data management in heterogeneous environments in which cooperation between networks and applications is advocated [11]. The last paper proposed techniques for incorporating travel-speed prediction in moving object databases [16]. This paper as the second paper is based on the concept of trajectories.

Session 3: The third and final paper session of the day focused on *Consistency and Replication* and included two papers in the emerging area of wireless data access for sensor networks. This session included four long talks and 1 short talk. The first talk presented a new event-based communication model for wireless multi-hop networks of energy-constrained devices such as sensor networks [4]. The network is arranged as an event dissemination tree, with nodes subscribing to the event types they are interested in. The next paper described a general purpose distributed index lookup service for mobile devices [7], which stores entries in form of (key, value) pairs in index caches located in each mobile device. Index caches are filled by epidemic dissemination of popular index entries. This was followed by the second paper on sensor networks, a paper on a scheme called TiNA. TiNA attempts to minimize energy consumption when performing in-network aggregation in a sensor network by exploiting data quality allowances specified by users [13]. Preliminary results show that TiNA can reduce power consumption by up to 50% without any loss in the quality of data. The final long paper was on media replication techniques in wireless peer-to-peer networks [5], and described a novel streaming architecture consisting of home-to-home online (H2O) devices that collaborate to provide on-demand access to a large selection of audio and video clips. The last paper was the short paper that presented algorithms to merge and reconcile XML data that is broadcast to mobile devices [8]. These algorithms were implemented in a tool called 3dm and used in two example applications, namely, a shared photo library and directory tree synchronization in a mobile file system. This session triggered a lively discussion along the topics of sensor networks and mobile peer to peer networks and disconnected operations.

The day concluded with an open floor discussion on mobility and pervasiveness, which build on the discussion from the last paper session. The central theme was the relationship (similarities and differences) between data management for mobile and tiny devices (e.g., sensor networks) and for peer to peer networks. The closing discussion also included thoughts on the format and frequency

of future MobiDE workshops. At this time, the participants overwhelmingly voted for the current format and co-location with ACM SIGMOD and ACM SIGMOBILE on alternate years. The vote was split regarding the frequency of holding the workshops – half of the participants voted to hold the workshop every year, and half voted to hold the workshop every other year.

3 Conclusions

MobiDE 2003 was very successful. The high-quality talks and papers resulted in a lively and informative discussion that carried through the entire workshop. The proceedings of the workshop have been published by ACM [1] and the workshop website [9] has more information about the workshop and its organization. The next MobiDE, the 4th one in the series, will take place in conjunction with ACM SIGMOD 2005, in Baltimore, MD on June 12, 2005. More information on MobiDE 2005 can be found at <http://db.cs.pitt.edu/mobide05>.

4 Acknowledgments

MobiDE 2003 would not have been successful without the contributions of all the members of the organizing committee and the members of the technical program committee. MobiDE 2003 was sponsored by ACM SIGMOBILE in cooperation with ACM SIGMOD, and also received support from ABC Virtual, Intel Research, Microsoft Research, IBM Research, and Hewlett Packard Labs. As well, as stated previously, a grant from the National Science Foundation was used to support the travel grant competition for graduate students.

References

- [1] *Proceedings of the Third ACM International Workshop on Data Engineering for Wireless and Mobile Access, MobiDE 2003, September 19, 2003, San Diego, California, USA*. ACM, 2003.
- [2] Ahmad S. Al-Mogren and Margaret H. Dunham. Concurrency control performance in days. In *MobiDE* [1], pages 25–29.

- [3] Sujata Banarjee and Panos K. Chrysanthis. Special issue on mobile data management: Guest editors' preface. *ACM SIGMOBILE Mobile Computing and Communications Review*, 8(3):1–3, July 2004.
- [4] Ugur Çetintemel, Andrew Flinders, and Ye Sun. Power-efficient data dissemination in wireless sensor networks. In *MobiDE* [1], pages 1–8.
- [5] Shahram Ghandeharizadeh and Tooraj Helmi. An evaluation of alternative continuous media replication techniques in wireless peer-to-peer networks. In *MobiDE* [1], pages 77–84.
- [6] Roger Keays and Andry Rakotonirainy. Context-oriented programming. In *MobiDE* [1], pages 9–16.
- [7] Christoph Lindemann and Oliver P. Waldhorst. Consistency mechanisms for a distributed lookup service supporting mobile applications. In *MobiDE* [1], pages 61–68.
- [8] Tancred Lindholm. Xml three-way merge as a reconciliation engine for mobile data. In *MobiDE* [1], pages 93–97.
- [9] MobiDE 2003 Organizing Committee. MobiDE 2003 workshop website, September 2003. URL: <http://db.cs.pitt.edu/mobide03>.
- [10] Declan O'Sullivan and David Lewis. Semantically driven service interoperability for pervasive computing. In *MobiDE* [1], pages 17–24.
- [11] Calicrates Policroniades, Rajiv Chakravorty, and Pablo Vidales. A data repository for fine-grained adaptation in heterogeneous environments. In *MobiDE* [1], pages 51–55.
- [12] Maya Rodrig and Anthony LaMarca. Decentralized weighted voting for p2p data management. In *MobiDE* [1], pages 85–92.
- [13] Mohamed A. Sharaf, Jonathan Beaver, Alexandros Labrinidis, and Panos K. Chrysanthis. Tina: a scheme for temporal coherency-aware in-network aggregation. In *MobiDE* [1], pages 69–76.
- [14] Goce Trajcevski. Probabilistic range queries in moving objects databases with uncertainty. In *MobiDE* [1], pages 39–45.
- [15] Xiaoxin Wu, Biswanath Mukherjee, and Bharat K. Bhargava. A low-cost, low-delay location update/paging scheme in hierarchical cellular networks. In *MobiDE* [1], pages 46–50.
- [16] Bo Xu and Ouri Wolfson. Time-series prediction with applications to traffic and moving objects databases. In *MobiDE* [1], pages 56–60.
- [17] Jianting Zhang and Le Gruenwald. Efficient placement of geographical data over broadcast channel for spatial range query under quadratic cost model. In *MobiDE* [1], pages 30–37.