

Dagstuhl Seminar “Mobile Information Management” – Executive Summary

Margaret Dunham¹, Birgitta König-Ries², Evaggelia Pitoura³, Peter Reiher⁴,
and Can Türker⁵

¹ SMU, Department of Computer Science and Engineering, Dallas, Texas, USA,
mhd@engr.smu.edu

² Friedrich-Schiller-Universität Jena, Institute of Computer Science, 07743 Jena,
Germany, koenig@informatik.uni-jena.de

³ University of Ioannina, Computer Science Department, GR 45110 Ioannina,
Greece, pitoura@cs.uoi.gr

⁴ UCLA, Computer Science Department, Los Angeles, CA, USA,
reiher@cs.ucla.edu

⁵ ETH Zurich, Institute of Information Systems, CH-8092 Zurich, Switzerland,
tuerker@inf.ethz.ch

1 Introduction

From October, 25th until 29th, 2004, a Dagstuhl seminar on mobile information management took place. The seminar was attended by 34 researchers from Europe and North America with backgrounds ranging from database systems, mobile information systems, geographic information systems and business informatics to wireless networks and security.

The motivation for organizing this seminar was as follows: The combination of wireless and wired connectivity along with increasingly small and powerful mobile devices, such as laptops, personal digital assistants, handheld/tablet PCs, and smart phones, but also embedded devices and sensors, enables a wide range of new applications and new ways to use existing applications. This will radically change the way information is managed and processed today. Information becomes ubiquitous, highly distributed and at the same time accessible from everywhere at any time. Information access takes place in highly dynamic and instable networks. Nevertheless, users and application developers will expect information processing to continue under similar guarantees as those offered by today's stationary and more or less centralized systems, even if some nodes of the ubiquitous information network are (temporarily) disconnected and/or are in motion. Examples for these guarantees are the ones given by database management systems, i.e. consistency and durability of data. Additional challenges stem from the fact that mobile devices while evolving fast, are (and will always be) less powerful than their stationary counterparts; they are smaller (resulting, e.g., in smaller in- and output devices and less available storage), have restricted energy supplies and communicate via expensive, unreliable wireless communication media and require new solutions to major existing security problems.

The goal of the seminar was to bring together researchers, especially from the area of databases and information systems, to identify open problems and new challenges in data, service, and user management in mobile information processing environments. We were particularly interested to provide a forum for discussing the consequences of the mobility of users and devices on today's and future data management systems. The first aim of these discussions was to clearly identify these consequences. The second and more challenging aim was to determine where existing solutions can be applied, where mobility raises truly new challenges, and which of these challenges are there to last. Of course, these discussions were to take into consideration not only general solutions from the database community at large, but also the approaches that have been developed by the mobile databases community over the last decade, e.g.:

- Moving objects and mobile users
- Mobile data dissemination and delivery
- Mobile data replication and synchronization
- Discovery and composition of mobile services
- Mobility awareness and adaptability
- Location-dependent, context-based querying and optimization
- Designing location-aware, context-aware services
- Mobile sensor and stream data management
- Continuous querying
- Self-organizing, self-tuning mobile components
- Ad-hoc processes and networks
- Mobile transactional processing
- Quality of Service for mobile databases
- Low power cryptography and other cryptographic solutions designed for wireless networks
- Methods of identifying, authenticating, and safely integrating mobile devices into a network

The discussions in this seminar were supposed to result in the identification of promising avenues of research that should be pursued in order to address the new issues that are introduced by mobility.

In order to achieve the goals mentioned above and in particular to foster discussions, the seminar was organized along two main lines: First, a number of overview talks were given. A brief summary can be found in Section 2. Second, a considerable amount of time was spent working in smaller groups of five to fifteen participants. Here, a number of topics were looked into more detail, research challenges and open issues were identified, avenues for future research were looked for. A summary of the working group activities is contained in Section 3.

2 Overview Talks

Throughout the seminar, a number of overview talks were scheduled. Extended abstracts of these talks can be found elsewhere in this volume.

The aim of these overview talks was on the one hand to reach a common basic understanding of a particular area before taking a closer look at the area in a working group, on the other hand, a number of areas were covered that we deemed important to the field, but where most people working in the field actually lack a deeper understanding.

The talks given and the main points made were:

- **Franz Lehner, Key Pousttchi: A Business Informatics View on Mobile Data Management.** Mobile Applications are not as successful in the market place as could be expected. The main argument why this is the case seems to be lack of knowledge about market mechanisms and players in the development community. New mobile applications need to be developed in a way compatible with business models.
- **Rainer Malaka, Pedro Jose Marron: Ubiquitous Computing.** While there has been a tremendous amount of research in this area over the last few years and first applications are reaching the market place, there remain a number of challenging open issues to be solved. Most of these issues will need the cooperative effort from more than one research community. Example issues are security, business models, data management, batteries, context-awareness.
- **Peter Reiher: Security for Wireless Networks and Mobile Environments.** The talk showed just how vulnerable mobile and wireless systems are and emphasized the need to develop solutions with security in mind. It also became evident that security solutions for wired networks often are not usable in mobile and wireless systems. Thus, a number of research questions were identified.
- **Can Türker: Transactions and Synchronization in Mobile Environments.** The talk gave an overview of both commercial and research approaches to transactions and data synchronization. It showed the wide gap in sophistication between the two. Also, while research has been quite active on mobile transactions for a while, it seems to have concentrated on rather classic problems and approaches. As of yet, approaches to handle, e.g. transactions in self-organizing networks or the new issues brought by the introduction of service-oriented computing are still missing.
- **Evaggelia Pitoura: Mobile, Ad-hoc, Peer-to-Peer, Self Organizing Systems: Data Engineering Issues.** This talk mentioned something that also was quite evident at the seminar: A lot of the problems addressed in mobile data management reappear in peer-to-peer and self-organizing systems. There is reason to hope that the solutions developed in mobile data management will also carry over to these new areas. Also, taken together, these areas form the basis of future ubicomp applications.
- **Michael Klein: Languages for (Semantic) Service Description.** Service-Oriented is a promising paradigm for mobile computing. On the one hand, since mobile devices possess comparatively little storage capacity, computing power etc., they often rely on functionality (services) provided by other mobile or non-mobile devices. On the other hand, mobile users

often find themselves in unknown environments, where they need to discover needed functionality. If this is to happen transparently to the user, semantic service-descriptions are needed that expressive enough to allow for automatic service discovery and binding.

3 Working Groups

During the seminar, most of the work was done in working groups. We had two sets of three working groups each. Each group got together for about a day, working on a more (or in one case: less) specific topic. Working group sizes ranged from five to fifteen participants. Working group sessions were concluded by a presentation of the results to the plenum.

The following working groups met. Again, more detailed descriptions of the individual working groups can be found elsewhere in these proceedings:

WG1: Mobile Business. This working group tried to get a better understanding of the key players in mobile business, how powerful they are and what their relationships are. A clear understanding of this system is necessary to enable successful development and deployment of mobile applications.

The main results of this working group were:

- There is a strong need for cooperation between the m-business informatics and the mobile data management community in order to ensure success of mobile applications. As a result, the foundation of an appropriate working group within the German Informatics Society bringing together researchers and practitioners from both areas will be proposed.
- Second, while the working group came up with an initial market model, this clearly is too vague and based on too many unsubstantiated assumptions to be of much use. Clearly, a thorough examination of these issues is necessary. There is a strong need for future research in this area.

WG2: Mobile Transactions. The group looked at different scenarios in which mobile transactions are needed, ranging from infrastructure-based networks to ad-hoc networks. Despite the large amount of work that has been done on mobile transactions in the past, a number of challenging research questions remain to be addressed, in particular with respect to networks with little or no infrastructure:

- Optimization of protocols with respect to the limited resources in wireless networks
- Adaptation of protocols to handle situations where the participating nodes change dynamically and are not necessarily all known to each other.
- Strategies to maintain transaction guarantees while allowing for, e.g., cooperative caching
- Strategies for replication decisions in dynamic networks.

WG3: Mobile Queries. After classifying different kinds of mobile queries and identifying the different dimensions that make mobile query processing so challenging, the working group concentrated on the role of context to ease dealing with mobile queries. There, a number of open research questions were highlighted.

- How is the context relevant to an application identified?
- Can research in recommender systems be utilized in recognizing relevant context, i.e. can context be “recommended”?
- What is an appropriate context model? Is a “*SAP-based*” approach where a common framework is defined and customized by each application more appropriate for context than a “*google*” approach where there is no common schema but rather a set of keywords?

WG4: Mobile Application Design. This working group started with the question “Why is the development of mobile applications harder than that of “normal” software?” In an attempt to answer this question, a number of mobile application scenarios and common experiences with the development were collected. It became evident that mobile application design requires to take a high number of dimensions, most of which are somewhat interrelated into account simultaneously. The group drew up an initial list of these dimensions.

As a result of the work, the following future activities are planned:

- The group intends to organize a special issue of a journal on this topic.
- The group’s report together with contributions from the journal should be extended to a handbook for mobile application development.

WG5: MANETs, P2P and Self-Organization. This working group addressed a rather broad set of questions. After identifying similarities and differences between the three classes of systems regarded, the group concentrated on identifying open research issues in all three areas.

- While traditionally, layered architectures where one layer hides its internals from the others seemed a good idea, in these systems, a certain amount of cross-layer interaction seems inevitable to ensure satisfactory performance. However, it is unclear how this interaction can be achieved and to which extent it is necessary.
- Mechanisms for cooperative caching and storage are still lacking.
- In self-organizing systems, the main challenge is to find ways to ensure the desired overall behavior by specifying local rules. This is made even harder by the fact that it cannot be assumed that all participants will actually adhere to the rules.

WG6: Description and Matching of Services in Mobile Environments. The working group classified services in mobile environments along two dimensions: The mobility or non-mobility of the service provider and the location-dependence or independence of the service offered. They suggest that this classification should

be reflected in the service description. Therefore, they developed an initial idea how service descriptions should be structured to be suitable in mobile environments: Descriptions should be split into a static part (containing the regular service description) and a dynamic part (describing the current context of the service provider). Depending on the class of service, the dynamic part will be more or less complex and important. Analogously, service requests should be split into several parts. Again, depending on the class of offer and request, different algorithms for matching of offers and requests need to be used. While there was general agreement that the idea seemed plausible, its realization will require considerable research effort.

4 Results of the Seminar

Mobile information management is a topic that is of immense and growing interest to a number of communities. It seems important to bring these communities together to ensure the development of appropriate solutions. While the databases community certainly is strong in questions of mobile information management and has a lot of solutions to offer, input from other researchers is needed: For instance, security issues are neglected in a frightening way in today's mobile applications and in particular in research prototypes. If we want these developments to enter the mass market, this needs to change. Research is often carried out without knowledge about or focus on existing business models etc. This, too, obviously, hampers market success.

A second result of the seminar was the identification of a number of promising avenues of research. These were discussed in more detail in the working groups (see above or in the individual working group reports). *should we summarize these here once again?*

Mobile information management is clearly an area of growing importance, and the research community should focus greater attention on providing the kinds of systems and services required to allow the field to reach its potential for creating new industries and improving people's lives. Mobility is not just an add on, but brings with it challenges that are different enough from more traditional system to warrant dedicated research.

Also, research on mobile data management offers a solid foundation for work on peer-to-peer and self-organizing systems and ultimately for ubiquitous computing. Again, this indicates the importance of this research area.