# Research Issues in Mobile Transactions**

Joos-Hendrik Böse[1], Stefan Böttcher[2], Le Gruenwald[3], Evaggelia Pitoura[4], Peter Reiher[5], George Samaras[6], Thomas Schwarz[7], Can Türker[8]

[1] Freie Universität Berlin, Institut fuer Informatik, Takustr. 9, 14195 Berlin, Germany
boese@mi.fu-berlin.de
[2] University of Paderborn, Computer Science, Fürstenallee 11, 33102 Paderborn, Germany
stb@uni-paderborn.de
[3] The University of Oklahoma, School of Computer Science, 200 Felgar Street, Room 116 EL, Norman, OK 73019, USA
ggruenwald@ou.edu
[4] University of Ioannina, Computer Science Department, GR 45110, Ioannina, Greece
pitoura@cs.uoi.gr
[5] UCLA, 405 Hilgard Ave, Los Angeles, CA 90095, USA
reiher@cs.ucla.edu
[6] University of Cyprus, Computer Science Department, CY-1678 Nicosia, Cyprus
cssamara@cs.ucy.ac.cy
[7] University of Stuttgart, Institute of Parallel and Distributed Systems (IPVS), Universitätsstraße 38, 70569 Stuttgart, Germany
thomas.schwarz@informatik.uni-stuttgart.de
[8] ETH Zürich, Department of Computer Science, ETH Zentrum, IFW C 47.2, CH-8092 Zurich, Switzerland,
tuerker@inf.ethz.ch

**Abstract.** This document focuses on transaction processing on mobile clients accessing a common set of data. The key requirement is to maintain all transaction properties as far as possible. We discuss three scenarios for databases on mobile clients, summarize typical applications and requirements for each of the three scenarios, and outline open research issues within each of the three scenarios. While the first scenario consists of mobile clients that are connected to a wired network, the second scenario consists of a network of mobile clients with a single-hop distance to each other but without a wired network, and the third scenario considers a network of mobile clients some of which are in multi-hop distance. The focus of our contribution is on distinguishing open problems from known solutions for each of the three scenarios. As the scenarios are increasingly challenging, we have identified additional requirements that occur within each new scenario compared to the previously discussed scenario.

** This document summarizes the results of the working group discussions on "Mobile Transactions" held during the Seminar on Mobile Information Management in October 2004 at Dagstuhl, Germany.

Dagstuhl Seminar Proceedings 04441
Mobile Information Management
http://drops.dagstuhl.de/opus/volltexte/2005/168

# 1 Introduction

The increasing capabilities of mobile devices in terms of processing power and connectivity drive the current trend to increasingly use such devices for accessing and especially altering mission-critical data on the go. While this greatly increases the user's flexibility and speeds up interactions, it also introduces a whole bunch of new problems. When moving from fixed-wired database technology to architectures including mobile clients, the following aspects play a key role to mobile transaction management:

- data has to be exchanged using small bandwidth connections that have to be shared with other mobile clients,
- unforeseeable disconnections and lost connections to mobile clients must be handled,
- movement of clients requires mobile transaction support,
- energy consumption limits the ability to send and receive messages,
- transaction models have to support these new requirements,
- scalability becomes a challenging requirement.

In our opinion, these requirements should guide further research in the field of data management for mobile computing. In the following sections, we outline three important scenarios for mobile data management and how each one affects transaction management. While all three scenarios have some research issues in common, the second and the third scenarios are more challenging than the first one.

# 2 Mobile clients with single-hop distance to a wired network (Scenario 1)

### Characteristics

Scenario 1 consists of mobile client nodes that communicate via wireless connections with a database server(s) residing in a fixed-wired network. The mobile clients store and operate on replicas of part of the data of the database. Updates of the mobile clients have to be merged or integrated with the data stored on the database server. Possibly, the updates have to be propagated to other mobile clients as well. We do not discuss the case where the mobile client acts as a dumb terminal only utilizing wireless connectivity as these scenarios presents minimal challenges.

Mobile nodes are connected to the fixed network via base stations that provide connectivity for a certain location or cell. Mobile nodes can lose their connection to a base station. They can move from one cell into another, switching the base station that provides them connectivity. In most situations, the speed of the mobility is slow compared to the speed of the network.

### Applications
Applications in such a setting include field service dispatching (e.g. maintenance technician at customer sites), distributed planning (e.g. resource planning), mobile commerce transactions (e.g. complex interactive orders) and support for traveling salesmen who must keep a large amount of data on their mobile computer to use while working with a remote customer.

### Requirements
Mobile transaction management has to provide for transparent disconnections, offline processing of data, and handing off the transaction context when the client moves to another cell. The main problem in this scenario is synchronizing the data updates performed locally at each mobile client with the data at the database server, possibly at a later point in time. To minimize conflicting changes, updates have to be propagated quickly to other affected mobile clients. A key requirement is scalability of transaction protocols to thousands of mobile clients.

### Open Issues
Although a lot of research has been done to address this scenario's problems, there is still demand for better scalability of protocols to support more frequent disconnections and larger numbers of clients and servers. In our opinion, the following are the open research issues for mobile transactions within Scenario 1:

- Performance optimizations with respect to limited energy resources on the client and with respect to small bandwidth connections between the server and the client.
- Scalability of protocols with respect to frequent disconnections, feasible for 100s -1000s of mobile clients per cell that have to be managed and synchronized by the server in the fixed-wired network.
- Handling unpredictable movements and unpredictable disconnections of clients.

## 3 Mobile clients with single-hop distance to each other but without a wired network (Scenario 2)

### Characteristics
The key characteristic of Scenario 2, which is its main difference from Scenario 1, is the absence of both a stable database server and a reliable fixed-wired network that the mobile clients can communicate with. Therefore, all communication relies on an unstable radio network, and the data of each particular node in the network may be lost in case of node failures. Another characteristic that distinguishes Scenario 2 from Scenario 3 is that within Scenario 2 all participating mobile clients are in single-hop distance to each other, and we do not anticipate network partitions. A typical case would be a group of participants that wants to exchange and process data between their nodes with transactional guarantees, given that participants may join and leave the group at unpredictable times.

## Applications

Applications that follow Scenario 2 include small groups that cooperate in a café or in a train. Further applications are mobile game parties, and bidding and betting in a mobile scenario without a fixed-wired network. Additional applications deal with emergency issues, such as fire protection, storm or flood management - as far as a single-hop connection between all participants can be guaranteed.

## Additional requirements (compared to Scenario 1)

Besides the requirements summarized for Scenario 1, the following requirements have to be met for Scenario 2. The system must be able to manage an unpredictably changing set of participants. Furthermore, transactional properties shall be granted as far as possible in the absence of a fixed-wired network and a stable database server.

## Additional open issues (compared to Scenario 1)

In addition to the open issues mentioned for Scenario 1, the following research questions should be considered.

- A first group of research issues covers how to guarantee transactional properties (isolation, atomicity and persistency) without a stable central coordinator. Research on preserving persistency includes techniques that avoid or limit information loss in case of a failing participant. Research on transaction concurrency control includes conflict definition, isolation and compensation models.

- A second group of research topics are related strategies for cooperative caching and data replication. Research issues on cooperative caching include how to guarantee cache consistency without a fixed server. Related to the problem of cache consistency is the handling of replicated data when the node holding the master copies disappears.

- A third group of research topics focuses on optimization issues and on the use of limited resources, e.g. bandwidth and energy. This is related to all other aspects of mobile transactions - ranging from query processing and caching to transactional properties (e.g. persistency, atomicity and isolation).

- A fourth group of research topics relates to the security of this scenario. The mobile users may have differing levels of trust in each other, and thus may need to control which users are able to access which data in the ad hoc network. Controlling data access in this environment requires some form of authentication, a method of expressing and reasoning about trust between participants, and interactions with other data management mechanisms, such as access control and the cooperative caching system.

## 4 A network of mobile clients with multi-hop distance (Scenario 3)

### Characteristics

In Scenario 3 (as in Scenario 2), we neither have a fixed-wired network nor a stable database server. But Scenario 3 is more difficult because multi-hop wireless routing is required to maintain communications between at least some of the mobile clients some of the time. In most situations where this scenario occurs, the clients providing the multi-hop routing are themselves moving, and often are subject to battery power constraints. As a consequence of the multi-hop network, we assume that the network of mobile clients may become partitioned, and from time to time the reachability of some participants is unknown. This generally complicates transaction processing.

### Applications

Applications that follow Scenario 3 include mobile workers in emergency scenarios, like fire protection, storm, earthquake and flood management - as far as a multi-hop connection between at least some participants is required, e.g. because of long distance communication or limited power supply. Another large group of applications is centered around car-based mobile information systems, including but not limited to traffic management or mobile commerce for car users. A further group of applications are military applications that cannot use a fixed-wired network. A final group of applications includes distributed games, i.e. games that do not rely on a direct connection to a server and that use multi-hop communication. If the many difficult underlying problems associated with these environments are eventually well solved, the convenience of being able to provide networking without setting up infrastructure ahead of time might significantly expand the scope of this scenario.

### Additional requirements (compared to Scenario 2)

Caching and replication strategies for multi-hop data management must consider network partitioning. The system should be operable during network partitions and a rejoining of different partitions should happen in a transparent way for the individual node. Furthermore, transaction management is more difficult to preserve because of an unknown reachability of participants. Also, energy limitations on all devices make it even harder to perform mobile transactions. Depending on the exact circumstances in which the ad hoc wireless network operates, individual participants might choose to act in manners that maximize their benefits at the cost of overall network performance (such as not spending their own battery power to forward other nodes' messages, even if the routing algorithm says they should), which would further complicate algorithms that handle mobile transactions and other data management issues. There are also serious security issues related to this scenario that tend not to arise in the earlier scenarios since mobile ad hoc networking often relies on the kindness of strangers.

### Additional open issues (compared to Scenario 2)

Beyond the research issues that have to be solved in Scenario 2, we consider the following issues to be important.

- One research issue is which transactional properties (i.e. atomicity and isolation) can be guaranteed up to which degree when operating in partitioned mode, i.e. without visibility of all participants.

- Another issue is related to caching strategies that consider multi-hop data management. In cases where not all participating nodes are equally trusted, either caching locations must be chosen with trustworthiness in mind, or data must be stored in an encrypted form.

- Energy issues become more important.

- Further problems are related to a layered approach to networks that exclude useful information (e.g. neighborhood of a node, costs of routes etc.) from the application, i.e. the database search and query processing. In our opinion, cross-layer approaches that allow for an inter-layer cooperation (e.g. between routing and database management operations) may be a promising starting point for a variety of new optimizations.

- Except in scenarios where all participating nodes are defined to be completely trustworthy and cooperative, the data management system cannot fully depend on the network to provide even basic behaviors like best-efforts message delivery. To some degree, these problems are akin to the partitioning and reachability issues arising from mobility and node failures, but active lying and deception by some nodes can lead to failures that are Byzantine, not simple. Handling such problems is more complicated.

## 5 Summary and conclusions

We have summarized typical applications and have discussed the requirements and open research issues of three increasingly challenging scenarios for mobile transactions. In order to identify common issues in all three scenarios, we have focused on the increasing challenges, while moving from one scenario to the next. One additional important research issue that has not been mentioned yet will be the theoretical limitations and unsolvable problems involved in the processing of mobile transactions within these scenarios.

Another question is the priority of the various research outlined here. Generally, the remaining issues related to scenario 1 should probably have a high research priority. This scenario is rapidly becoming a reality, and the scaling issues that remain unaddressed are likely to be a problem in upcoming years, if they are not addressed relatively soon. The issues for the scenarios 2 and 3 are probably of lower research priority, though they are also important and need the researchers' attention. While not a common mode of operation for data management today, increasingly ubiquitous computing environments and the growth in the use of mobile devices will likely require data management systems to handle the situations described in Scenario 2 and in Scenario 3 in the not-distant future. Furthermore, car-based mobile applications may become a driving force for research in Scenario 3. Thus, the time to address these research issues is now. As underlying networking mechanisms to handle multi-hop mobile ad hoc networks effectively are in their infancy, we are unlikely to see large deployments of data management systems on top of a multi-hop network in the near future. However, very likely data management systems that do eventually operate in this environment will require tighter interactions with the networking layer than is required for other scenarios. If we delay research on data management issues in this area until the networking community has produced mature technologies for ad hoc mobile wireless environments, it may be too late to achieve those tighter interactions. Thus, while most applications require research efforts in data management for mobility according to the first scenario, it is vital that some significant efforts be put into the more far-reaching and challenging problems of the second scenario and the third scenario.