Mobile Database Research: What Is To Be Done?

Le Gruenwald
University of Oklahoma
School of Computer Science
ggruenwald@ou.edu

Frank Olken
Laurence Berkeley National Laboratory
and National Science Foundation
folken@nsf.gov

1. Current State of Mobile Database Research

Work on mobile databases started in the early 90's. Much research in this area assumes a cellular network where there is a fixed Mobile Support Station (MSS or server) that supports all mobile hosts (MHs or clients) roaming within its cell. When an MH moves out of a cell and enters a new cell, it can no longer communicate with the previous cell's MSS, and is under the control of the new cell's MSS. All MSSs communicate with each other via a fixed network. So in many of the existing works, mobile database systems refer to DBMS client-server systems in which the client is mobile (e.g., runs on a laptop, PDA, or smart cellular phone), the server is static, the client has intermittent connectivity, the client self-identifies its location (e.g., via GPS), and the queries are often "location aware" (e.g., find me the nearest Chinese restaurant). Other common applications include vehicle tracking (military, emergency vehicles, taxis) where some (or all of the clients) are also mobile. The key research issues that the research community has tried to tackle are how to cope with mobility, frequent disconnection and energy limitation of the client in a number of DBMS issues, such as transaction management, data caching, data replication and location-aware query processing [Serrano-Alvarado, 2004]. Additional work is also found in the areas of security and privacy [Bernard, 2004], especially for detecting malicious transactions and preserving privacy in location-aware queries.

Following cellular networks is another network structure, called Mobile Ad-hoc Network (MANET). In this network, all MSs (clients) are roaming and the network that interconnects these MHs is a wireless network with a frequently changing topology, and there are no fixed infrastructures or fixed MSSs (servers). Some research efforts attempting to handle database issues in MANET have slowly emerged ([Brayer, 2005] [Padnabhan, 2008]).

Besides the client-server architecture, the database research community also examines the P2P architecture ([Brayer, 2005], [Luo, 2008]) where peers communicate with each other without going through a central node.

However, the database issues are far from being resolved for both cellular networks and ad-hoc networks when one considers dynamic applications with continuum domain data, such as AWACS radar and storms tracking, and sensing applications in sensor wireless networks. In many current works, the most popular applications are location-aware queries for mobile clients and static targets (stores), and navigational queries (shortest path, fastest path) over static (or slowly changing) road networks, where typically locations of vehicles are known via self-reporting (GPS). Some works on mobile database systems are not even concerned with location-aware queries but rather with the issues of coping with mobile, intermittently connected client systems. In Section 2, this paper articulates some open research topics for mobile database research.

2. Open Research Topics

It is commonplace today for mobile database systems to consider the space in which the vehicles are mobile to be either a directed graph (road network) or a two dimensional metric space (i.e., a planar representation used in geographic information systems). Military and emergency operations applications will require 2.5 dimensional contexts – e.g., representation of the surface topography as constraints on vehicle mobility and affecting transit times. Systems involving control of aircraft or submarines will need three dimensional spatial representations in which the aircraft or submarines operate. Systems which concern themselves in the possibility of intersecting vehicle trajectories (usually trying to avoid collisions) may need to operate in 4 dimensions (3 spatial dimensions and one temporal).

Typical mobile database systems today treat “road delays (transit times)” as either static or exogenous time-varying variables to be observed and incorporated into trip planning. We need systems which can fully address time-dependent road delays (e.g., anticipate rush hour traffic patterns). We will also need systems which treat “road delays” as endogenous variables which can be influenced by vehicle routing decisions. Such considerations arise both in routine traffic control during rush hour, in the presence of accidents and other road closures, and for evacuation planning and control during floods, hurricanes, and large scale fires.

There has been relatively little research thus far on mobile database systems in which not only are the vehicles moving, but also the environment is mobile. Of particular note are systems involving tracking (or planning) vehicle movements in continuous domains (at sea, in the air, cross country motion of SUVs (or tanks), etc.) in which the environment is also mobile. Examples include moving weather systems (hurricanes, thunderstorms, etc) which must be avoided, or fast moving wild fires or forest fires. In these settings both the vehicles and the environment (weather, fire, ocean currents) are mobile. Applications include both routine air or sea navigation (where one must route aircraft or ships around severe weather systems) and emergency operations during large scale storms, hurricanes, floods, or fires. Routine maritime route planning considers ocean currents to minimize fuel expenditures and travel times. Such systems will effectively need to integrate mobile databases with spatio-temporal databases. Note that most of these applications involve frequent updates of the environmental descriptions (weather, fire, ocean currents).

Today most mobile databases assume that data values (vehicle locations, road segment transit times, road network connectivity) are known with certainty. We need systems which can cope with uncertain data concerning vehicle locations, road delays, connectivity, location and direction of weather systems (e.g., hurricanes).

Traditional (relational) database systems have difficulty integrating non-trivial computations with data retrieval. Similar problems arise in mobile databases involving routing computations, weather forecasting, etc. Open questions include how much computation (routing, traffic/weather forecasting) to include within the DBMS and what should remain outside the DBMS.

With the advancements of sensor networks, mobile sensing applications are possible. For example sensors attaching to a patient can send his health condition (e.g. heart rate) periodically to a hospital for health monitoring. Data stream issues, such as varying data distribution, frequent data missing and corruption, and infinite amounts of data, need to be considered together with mobile database issues. For example a query model for mobile sensor databases should be not only location-aware and energy-aware, but also data stream-aware.

Another issue is about energy consumption. Because of the energy limitation of battery-powered clients, the mobile database community has been the first database community that addressed the energy consumption issue. As GREEN IT has increasingly played an important role in research directions for both the academia and industry, it is important for the mobile database research community to extend its work on energy consumption to fixed database structures, which has originally been considered to have unlimited energy, in order to achieve a GREEN DBMS.

Last but not least is the topic of performance evaluation. Like in many database research areas, works in mobile databases often use synthetic data or a small set of real data to evaluate their performance. Justifications for the validity of such data are usually either not provided or provided with little detail. Scalability studies are usually ignored. For performance evaluation to be credible, it is important for the research community to develop a comprehensive repository of data and queries generated from real-world applications and a benchmark.

3. Conclusions

Increasingly we live in a world in which vehicles (autos, aircraft, ships) and persons will each carry location-aware (GPS) computing platforms which communicate wirelessly. The communications will still be intermittent. We will need to develop mobile database systems which can cope with millions (or billions) of such mobile clients and which cope with mobility in road networks, 2D, 2.5D, 3D settings, and which can cope with uncertain data, and mobility of both vehicles (persons) and the environment (weather, ocean currents, fires). We will need to also extend the systems to accommodate sensing applications where data streams characteristics are taken into account.

Acknowledgements

Le Gruenwald's work has been partially supported by NSF Award IIS-312746 on Mobile Database Systems. Frank Olken's work has been supported by the NSF IR/D program via an IPA grant from NSF to LBNL (Lawrence Berkeley National Laboratory).

References

[Bernard, 2004] G. Bernard, C. Roncancio, P. Serrano-Alvarado and P. Valduriez, "Mobile Databases: a Selection of Open Issues and Research Directions," *ACM SIGMOD Record*, vol. 33, No. 2, pp 78-83, 2004.

[Brayner, 2005] A. Brayner and F. S. Alencar, "A Semantic-serializability Based Fully-Distributed Concurrency Control Mechanism for Mobile Multi-database Systems," *Proceeding of the 16th International Workshop on Database and Expert Systems Applications*, pp. 1085-1089, 2005.

[Luo, 2008] Y. Luo and O. Wolfson, "Mobile P2P Databases," *Encyclopedia of GIS*, pp. 671-677, 2008.

[Padnabhan, 2008] P. Padnabhan, L. Gruenwald, A. Vallur, M. Atiquzzaman, "A Survey of Data Replication Techniques for Mobile Ad-Hoc Network Databases," *Journal of Very Large Data Bases*, Vol. 17, Issue 5, August 2008.

[Serrano-Alvarado, 2004] P. Serrano-Alvarado, C. Roncancio and M. Adiba, "A Survey of Mobile Transactions," *Distributed and Parallel Databases*, vol. 16, no. 2, pp. 193-230, September 2004.