

Location-Based Services—A Database Perspective

Christian S. Jensen, Anders Friis-Christensen, Torben B. Pedersen, Dieter Pfoser,
Simonas Šaltenis, and Nectaria Tryfona

Department of Computer Science, Aalborg University
Fredrik Bajers Vej 7E, DK-9220 Aalborg Øst, Denmark
`www.cs.auc.dk/~{csj|afc|tbp|pfoser|simas|tryfona}`

Abstract. We are heading rapidly towards a global computing and information infrastructure that will contain billions of wirelessly connected devices, many of which will offer so-called location-based services to their mobile users always and everywhere. Indeed, users will soon take ubiquitous wireless access to information and services for granted. This scenario is made possible by the rapid advances in the underlying hardware technologies, which continue to follow variants of Moore's Law.

Examples of location-based services include tracking, way-finding, traffic management, safety-related services, and mixed-reality games, to name but a few. This paper outlines a general usage scenario for location-based services, in which data warehousing plays a central role, and it describes central challenges to be met by the involved software technologies in order for them to reach their full potential for usage in location-based services.

1 Introduction

In the coming years, wireless computing will find widespread application, and computing will become increasingly ubiquitous. The wireless revolution promises to have as fundamental an impact on society as the emergence of the commercial Internet.

Some industry observers predict that there will be billions of wireless “appliances” on the Internet within relatively few years. These appliances include hand-held devices, e.g. mobile-phone terminals and the diverse range of PDAs already available, wrist watches, smart clothing and jewelry, cameras, home appliances, cars, smart dust, etc. This proliferation of on-line appliances is made possible by continued advances in electronics, including miniaturization, lower power consumption, broader band wireless connections, better display technologies, across-the-board sustained improvement in the performance/price ratio, etc.

Technologies such as WAP and Bluetooth and their successors will bring about the required connectivity. Positioning systems based on satellite technologies, currently the Global Positioning System (GPS), as well as technologies embedded in the wireless connectivity infrastructures, will enable the positioning of a substantial portion of the appliances and their users.

In the near future, the Internet will be accessed more frequently from wireless devices than from desktop computers. Some experts pertain that we have yet to come to understand how the above technologies will be used and how they will impact our lives.

However, near-term applications include position-enabled tourist services, safety and security services, metered services (such as car insurance), traffic services and management, etc. These applications will involve the management of vast quantities of data.

The scenario sketched here will see the applications of techniques from a wide range of disciplines in computer science. In particular, it will become a huge area of application for “geographic” as well as database management technologies. In this scenario, where change is inherently continuous, the capture, modeling, maintenance, and querying of location information is of critical importance. This paper describes some of the challenges that are raised by the need for offering location-based services to online, position-aware objects capable of continuous movement.

The next section introduces location-based services, and Sect. 3 covers fundamental data warehousing concepts in relations to location-based services. Sect. 4 then describes six of the challenges related to location management posed by location-based services. Sect. 5 summarizes the paper.

2 Location-Based Services

Following a description of the kinds of mobile objects that location-based services concern, this section offers a database-centric usage scenario for location-based services and exemplifies such services.

Online, Position-Aware, Wireless Mobile Objects. The coming years will witness very large quantities of on-line (i.e., Internet-worked), position-aware, wireless objects capable of movement. Examples of such objects include the following.

- Consumers using Internet-enabled mobile (phone) terminals (e.g., WAP, I-mode, WAP2.0) with enhanced displays, as well as diverse types of personal digital assistants (PDAs, e.g., PalmPilots). As examples, tourists may carry on-line and position-aware “cameras” and “wrist watches” will evolve to become on-line and position-aware, as will jewelry, “clothing,” and luggage.
- Vehicles with computing equipment, including automobiles, public transportation vehicles, recreational vehicles, sea vessels, etc. Luxury cars already carry navigation equipment, and a wide range of such equipment is available to consumers for integration into older vehicles as well as new economy vehicles. The major car manufacturers are working intensely with, e.g., the major mobile phone providers to integrate Internet access into their future automobiles.
- Home appliances. Using local-area wireless communications technologies (e.g., Bluetooth) within homes and wireline technologies beyond the homes, home appliances will increasingly be on-line, at lower and lower cost.

Location-Based Services Scenario. The following general, database-centric scenario is considered.

Moving objects use e-services that involve location information. The objects disclose their positional information (position, speed, velocity, etc.) to the services, which in turn use this and other information to provide specific functionality.

The services maintain a log of the requests made to them, and use this for analyzing user interaction with the services. In an Internet context, a sequence of requests is called a click-stream. The services accumulate data derived from the click-streams and integrate this with other customer data in data warehouses, which are very large repositories of integrated information used for data analysis (described further in the next section).

The data in the data warehouse is used for mass-customization of the services, so that each user receives a service customized to the user's specific situation, preferences, and needs. This involves the immediate generation of dynamic web-page content based on the warehouse and the user's current interaction with the service, e.g., the display of advertisements that are relevant to this particular user at the particular place and time.

In addition, the warehouse is used for delayed modification of the services provided, and for longer-term strategic decision making. Business intelligence techniques such as on-line analytical processing and data mining are used for these purposes.

The integration of location information into this scenario has as of yet received little attention and offers a number of fundamental challenges. Common to these challenges is the task of extending techniques that work well for static data to support the kinds of dynamic, continuously evolving positional data that is found in location-based services.

Example Location-Based Services. The five categories described next characterize what may be thought of as standard location-based services; they do not attempt to describe the diversity of services possible.

i *Traffic coordination and management.*

Based on past and up-to-date positional data on the subscribers to a service, the service may identify traffic jams and determine the currently fastest route between two positions, it may give estimates and accurate error bounds for the total travel time, and it may suggest updated routes for the remaining travel. It also becomes possible to automatically charge fees for the use of infrastructure such as highways or bridges (termed road-pricing and metered services).

ii *Location-aware advertising and general content delivery.*

Users may receive sales information (or other content) based on their current locations when they indicate to the service that they are in "shopping-mode." Positional data is used together with an accumulated user profile to provide a better service, e.g., ads that are more relevant to the user.

iii *Integrated tourist services.*

This covers the advertising of the available options for various tourist services, including all relevant information about these services and options. Services may include *over-night accommodation* at camp grounds, hostels, and hotels; *transportation* via train, bus, taxi, or ferry; *cultural events*, including exhibitions, concerts, etc. For example, this latter kind of service may cover opening-hour information, availability information, travel directions, directions to empty parking, and ticketing. It is also possible to give guided tours to tourists, e.g., that carry on-line "cameras."

iv *Safety-related services.*

It is possible to monitor tourists traveling in dangerous terrain, and then react to emergencies (e.g., skiing or sailing accidents); it is possible to offer senile senior citizens more freedom of movement; and it is possible to offer a service that takes

traffic conditions into account to guide users to desired destinations along safe paths.

v *Location-based games and entertainment.*

One example of this is treasure hunting, where the participants compete in recovering a treasure. The treasure is virtual, but is associated with a physical location. By monitoring the positions of the participants, the system is able to determine when the treasure is found and by whom. In a variation of this example, the treasure is replaced by a “monster” with “vision,” “intelligence,” and the ability to move. Another example in this category is a location-based ICQ service.

3 Data Warehouses for Location-Based Services

A brief introduction to data warehousing is followed by an example of a specific data warehouse, which is then used for illustration in the next section.

3.1 Introduction to Data Warehousing

A data warehouse is a large data repository that integrates data from several sources into structures expressly designed for *analytical* purposes. Data warehouses typically employ a *multidimensional model* for organizing data [4]. This type of model typically categorizes data as either business *facts* with associated *measures*, which are numerical in nature, or *dimensions*, which characterize the facts and are mostly textual.

For example, in a retail business, *products* are sold to *customers* at certain *times* in certain *amounts* at certain *prices*. A typical fact would be a *purchase*. Typical measures would be the *amount* and *price* of the purchase. Typical dimensions would be the *location* of the purchase, the type of *product* being purchased, and the *time* of the purchase.

Each dimension is organized into a hierarchical structure of *levels*, which enables the aggregation of facts to the desired levels of granularity. For example, the Time dimension may have levels *Day*, *Month*, *Quarter*, and *Year*. A prototypical query applies an aggregate function, such as SUM, to the facts characterized by specific values from the dimensions, e.g., to obtain the totals of sales by Month, Product Group, and State.

3.2 Data Warehouse Case Study

We proceed to describe a small case study of a data warehouse for LBS, which stores *requests* made to a service and makes them available for analyses. A diagram describing the schema of the warehouse is seen in Fig. 1.

Two of the five dimensions of this data warehouse concern time. The Date dimension captures the date the request was made. It has the *levels* Day, Week, Month, Quarter, and Year, as well as a \top level that represents *all* dates. The hierarchy illustrates how the different levels *roll up* to one-another, e.g., days roll up to weeks and months, but weeks do not roll up into months. The Time Of Day dimension captures the time during the day the request was made. It has been separated from the Date dimension to allow for analyses regarding the time of day across several days.

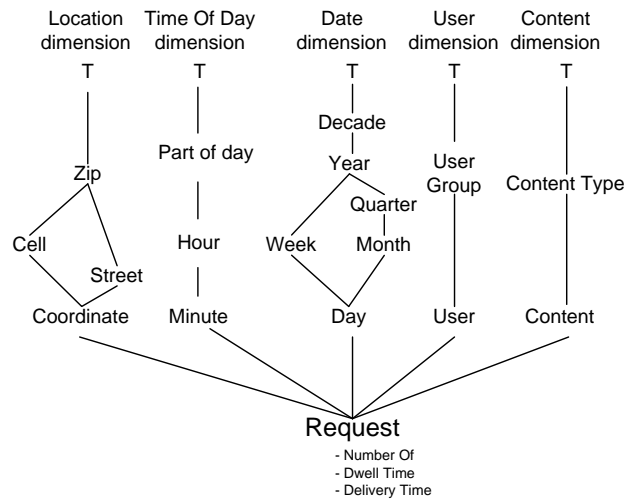


Fig. 1. Schema of Data Warehouse Case Study

The Content dimension of the warehouse captures the content of the service being requested, e.g., sports news or concert information. The precise type of content, e.g., Sports News is grouped into Content Types, e.g., News. The User dimension captures which user made the request, and additionally groups users into User Groups. The Location dimension captures the location from which a service was requested. Some user locations can be obtained very precisely, e.g., via GPS, while others can only be obtained at the level of the *cells* in the cellular transmission network.

The warehouse has three measures. The Number Of measure captures the number of requests made for a particular combination of dimension values. The Dwell Time records the number of seconds the user looked at the response to the request before making a new request. The Delivery Time measure captures the number of seconds used to serve the requests for the given combination of dimension values.

The warehouse is used for answering questions about the connection between the user, the user's location, the time of day, and the content of the service being requested. This can be used for providing customized services, e.g., custom content specific to the interest of the user, the user's location, and the date and time of day.

4 Research Challenges in Location-Based Services

We proceed to describe six research challenges posed by location-based services.

4.1 Support for Non-Standard Dimension Hierarchies

In traditional data warehouse environments, the dimension hierarchies must have a very regular structure; more precisely, they must be *balanced trees*. In contrast, location-

based services require support for *irregular* dimensions. As an example, consider the instances of the Location dimension in Fig. 2.

This hierarchy exhibits two irregularities. First, the hierarchy is *non-strict* [8], as cell 1001 is a child of more than one Zip value, namely the two Zip's 9000 and 9220. Second, the hierarchy is *non-onto* [8], as ZIP 9990 has no cells in it (this area is a wilderness areas without cellular network coverage).

Irregular hierarchies cannot be captured well by the standard multidimensional data models used in data warehouses currently, meaning that designers of an LBS warehouse cannot specify the dimensions they desire, and it opens to the problem of incorrect aggregation of data (double-counting or not counting certain values) [8]. In addition, existing optimization techniques that are essential to providing adequate query performance in very large data warehouses require regular hierarchies.

It follows that there is a great need for both data modeling and implementation techniques that support irregular dimension hierarchies. Only little previous research has addressed the issue of modeling irregular hierarchies [8, 9] and implementing systems with such hierarchies efficiently [7], and it remains to be seen whether this initial research is effective in an LBS setting.

4.2 Support for Imprecision and Varying Precision

Imprecision is a fundamental aspect of location data. User locations are sampled according to some specific protocol. As examples, (i) the users may disclose their locations when they desire a service, (ii) they may supply their positions at regular intervals, or (iii) they may keep track of where the service thinks they are and then issue position samples to the service exactly when necessary in order for the service's record to be within a certain threshold of the actual position. In this latter example, the threshold is dependent on the specific service desired. As a result of the sampling, complete traces of the users' movements are unavailable; rather, the service only knows the locations of the users at discrete times.

Additionally, the samples themselves are imprecise [10]. The sample imprecision is dependent on the positioning technology used and the circumstances under which a specific technology is used. For example, the cellular infrastructure itself, the positioning technologies offered by companies such as SnapTrack and Cambridge Positioning Technologies, and GPS and server-assisted GPS offer quite different precisions. And, for example, GPS technology is dependent on lines of sight to several satellites, which affects the robustness of the technology. In other words, the accuracy of the positioning is highly dependent on the user's location.

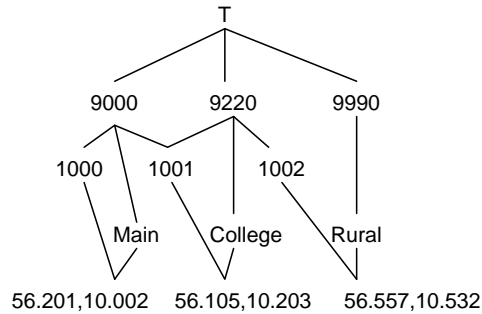


Fig. 2. Instance of the Location Dimension

In the case study, the locations captured by the Location dimension are imprecise, and this imprecision must be taken into account in the data representation and must be handled intelligently during querying. If a precise (but incorrect) trace is maintained for each user, the users querying the warehouse may make suboptimal decisions based on wrong positional information. On the other hand, if an overly imprecise record of the positions is kept, the warehouse also offers suboptimal support for queries.

A related challenge is that location data often have *varying precisions*. In the case study, some locations are obtained using GPS and can thus (often) be obtained at the quite precise level of Coordinate (within 10 meters). However, for a user without a GPS receiver, the location can only be obtained using the cellular network, where the most precise location is approximately that of a cell; this may lead to positioning that is imprecise in the range of hundreds of meters or more. It should be possible to represent position data with such varying precisions, and queries have to give meaningful answers, even though the data has this highly varying level of precision.

There has been quite a lot of research in the management of imprecise data in general databases [2], as well as some initial research on the handling of imprecision and varying precision in multidimensional data models [6, 8]. Additional research is needed to support the complexities of LBS, e.g., the handling of imprecision and varying precision in queries about complete *user traces* rather than single points.

4.3 Support for Movement Constraints and Transportation Networks

The movements of the users of location-based services are often subject to two types of constraints. In the *blocking objects* type of constraints, objects block the movement of users. As examples, the movements of skiers are blocked by terrain without snow, and the movements of sea vessels are blocked by land. In the *network* type of constraint, the movement of objects occurs in, and is constrained by, networks. Such networks encompass transportation networks, including road and rail networks, and the infrastructures of buildings consisting of rooms connected by walkways and stairways.

Knowledge of such constraints may be exploited by a service to better estimate the position of a user, which may lead to more accurate position information or a lower sampling rate.

Transportation networks¹ are particularly interesting. When a user's movement is constrained to such a network, the movement is effectively constrained to a space with a lower dimensionality [5]. For example, if a user is moving on a road and is interested in advertisements of sales, the problem of finding the locations of sales nearest to the user may in a sense be reduced to a one-dimensional problem, even if the user is really moving in what is perceived as a two-dimensional world.

In the case study, we would maintain a representation of the infrastructure, in which movement occurs, separately from the data warehouse proper. And we would attempt to determine the changing transportation modes of the users. For example, a user may be walking, using public transportation, or driving a car. The user may indicate the mode to

¹ Note that a transportation network is different from the mathematical notion of a (directed) graph. In a transportation network the (geographical) location of the nodes is significant, while the notion of position of the nodes in a graph is non-existent.

the service, or the service may determine the mode based on the user's location, speed, or general movement pattern, or based on yet other attributes.

4.4 Support for Multi-Resolution Objects and Maps in Data Modeling

Location-based services call for the availability of multi-scale objects and maps. Consider the situation in which a user approaches a new city and is looking for a highly recommended restaurant. Initially, the user follows a route that leads in the direction of the restaurant, as indicated on a large-scale map. Then, as the user enters and navigates the city, a more detailed map is presented to the user. This map offers a more detailed image of not only of the geographic area, but also of the target object. For example, a service that indicates a gas station at a coarse level can also include a restaurant and a mini market at a more detailed level.

This scenario indicates that in order to provide users of location-based services with meaningful results, we need to handle maps and objects at varying levels of detail.

Current conceptual and logical database modeling techniques only allow us to relate descriptive data (e.g., services) to geographic data at one level of detail. The extension of popular and well-known models such as the Entity-Relationship and the Unified Modeling Language and the relational model to capture multi-resolution maps and objects is highly desirable.

4.5 Support for Spatial Data Mining on Vehicle Movement

Location-based services will frequently be used from vehicles. In step with the increasing volume of vehicles, the need for further knowledge about possible routes and traffic conditions, as well as the prediction of troublesome situations gains in prominence. The area of computer science dealing with the discovery of spatial knowledge is called spatial data mining. Typical scenarios in which spatial data mining methods are useful include the following:

- Extraction of knowledge about the movement of vehicles based on already existing data. Examples are: “busy routes,” “number of users moving towards the city center,” and “routes on which cars speed up.”
- Prediction of the times and locations of troublesome situations, such as traffic jams.
- Provision of alternatives, such as alternate routes with different characteristics (e.g., small probability of accidents, low toll, scenic, small variation in travel time).

For the realization of this idea, theories from the areas of general data mining [3], spatial data mining [12], and trajectory management [11] need to be combined. In relation to GIS, this should result in efficient support for forecasting and risk analysis based on the analysis of moving-object and static data, including real-time data and stored, past data. To support location-based services, such analyses must occur in real-time.

4.6 Support for Continuous Location Change in Query Processing Techniques

Two classes of implementation techniques are essential to ensuring adequate performance of querying, namely indexing and precomputation. We consider the challenges posed to indexing by the continuous change of user locations.

Indexing is a fundamental technique in data management, as it makes it possible to locate desired data items in very large databases efficiently, typically in time logarithmically proportional to the total number of data items. Because data access is increasingly becoming a bottleneck, indexing techniques are becoming increasingly important.

Traditional indexing techniques work only for static data, meaning that the indices have to be explicitly updated when changes occur to the data. For large, continuous datasets, e.g., those capturing the positions of moving users, the constant updating of the indices would require very large computing resources, rendering the use of indices either impractical or totally impossible. Or, alternatively, the large volumes of updates and the mechanisms that regulate the concurrent use of the indices would combine to block the querying of these structures, also rendering them useless.

It is a fundamental challenge how to obtain, if at all possible, the well-known and widely relied upon benefits of indexing when the data being indexed change continuously.

Two general approaches may be taken towards accommodating continuity. Techniques may be applied that (i) create less updates, or (ii) the existing techniques may be enhanced to support rapid, non-bulk update.

One example of the former is the representation of the movement of a user by a position and a velocity vector, instead of simply by a position. With the velocity vector available, the position needs only be updated when this vector changes, which generally leads to less updates. The use of a transportation network to reduce the need for updates (in Sect. 4.3) is another example.

Many techniques may play a part in supporting rapid updates. For example, buffer techniques [1] may be applied. Briefly, these remedy the inefficiencies of transferring blocks with little data between main memory and disk by buffering updates.

5 Summary

Location-based services is rapidly emerging as a prominent area of deployment of “geographic” and data management technologies. It is argued that data warehousing may play a prominent role in such services, but that the existing, promising technologies must meet new and tough challenges in order for them to reach their full potential in this area.

The paper initially characterizes location-based services by exemplifying the moving objects that location-based services concern, by describing the general scenario in which the services may exploit data warehouse technologies, and by giving examples of specific location-based services. Following an introduction to data warehousing and a case study, the remainder of the paper is devoted to the discussion of research challenges posed by location-based services.

The challenges concern various aspects of the support for “location.” They stem from the complex types of spatial regions in which objects move, from the limitations

inherent in the capture of the locations of moving objects, from the complex nature of the captured data, and from the rapid database change implied by the capture of continuous movement.

It is hoped that these challenges will spur researchers to formulate more concrete challenges, to which they can subsequently develop new solutions. It should also be emphasized that the challenges brought forward here reflect the backgrounds of the authors; other authors may emphasize other, equally valid challenges.

Acknowledgments

The authors would like to thank Christian Hage of Euman A/S for sharing his insights into location-based and e-services with us. This work was supported in part by the WIM network, funded by NorFA through grant 000389, by the Danish Technical Research Council through grant 9700780, and by a grant from the Nykredit Corporation.

References

1. L. Arge. The Buffer Tree: A New Technique for Optimal I/O-Algorithms (Extended Abstract). In *Proceedings of the Fourth International Workshop on Algorithms and Data Structures*, pp. 334–345, 1995.
2. C. E. Dyreson. A Bibliography on Uncertainty Management in Information Systems. In A. Motro and P. Smets (eds.), *Uncertainty Management in Information Systems—From Needs to Solutions*, Kluwer Academic Publishers, pp. 413–458, 1997.
3. M. Ester, H.-P. Kriegel, and J. Sander. Knowledge Discovery in Spatial Databases. In *Proceedings of the 23rd German Conference on Artificial Intelligence*, LNCS 1701, pp. 61–74, 1999.
4. R. Kimball and R. Merz. *The Data Webhouse Toolkit*. Wiley, 2000.
5. G. Kollios, D. Gunopulos, and V. J. Tsotras. On Indexing Mobile Objects In *Proceedings of the PODS Conference*, pp. 261–272, 1999.
6. T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. Supporting Imprecision in Multidimensional Databases Using Granularities. In *Proceedings of the Eleventh International Conference on Scientific and Statistical Database Management*, pp. 90–101, 1999.
7. T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. Extending Practical Pre-Aggregation in On-Line Analytical Processing. In *Proceedings of the Twentyfifth International Conference on Very Large Databases*, pp. 663–674, 1999.
8. T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. A Foundation for Capturing and Querying Complex Multidimensional Data. In *Information Systems*, special issue on data warehousing, 42 pages, to appear, 2001.
9. T. B. Pedersen and N. Tryfona. Pre-Aggregation in Spatial Data Warehouses. In *Proceedings of the Seventh International Symposium on Spatial and Temporal Databases*, to appear, 2001.
10. D. Pfoser and C. S. Jensen. Capturing the Uncertainty of Moving-Object Representations. In *Proceedings of the Sixth International Symposium on Spatial Databases*, pp. 111–132, 1999.
11. D. Pfoser, C. S. Jensen, and Y. Theodoridis. Novel Approaches in Query Processing for Moving Objects. In *Proceedings of the 26th International Conference on Very Large Databases*, pp. 395–406, 2000.
12. N. Tryfona. Constraint-Based Spatial Data Mining. In *International Journal of Computers, Environment and Urban Systems*, to appear.