TRAX—Real-World Tracking of Moving Objects

Christian S. Jensen Stardas Pakalnis
csj@cs.aau.dk stardas@cs.aau.dk
Department of Computer Science, Aalborg University, Denmark

ABSTRACT

A range of mobile services rely on knowing the current positions of populations of so-called moving objects. In the ideal setting, the positions of all objects are known always and exactly. While this is not possible in practice, it is possible to know each object's position with a certain guaranteed accuracy.

This paper presents the TRAX tracking system that supports several techniques capable of tracking the current positions of moving objects with guaranteed accuracies at low update and communication costs in real-world settings. The techniques are readily relevant for practical applications, but they also have implications for continued research. The tracking techniques offer a realistic setting for existing query processing techniques that assume that it is possible to always know the exact positions of moving objects. The techniques enable studies of trade-offs between querying and update, and the accuracy guarantees they offer may be exploited by query processing techniques to offer perfect recall.

1. INTRODUCTION

The continued advances in consumer electronics, wireless communication, and geo-positioning combine to enable new kinds of mobile services that rely on knowing the up-to-date geo-locations of entire populations of moving objects. Examples include traffic congestion monitoring and various kinds of fleet management (police cars, dangerous transports, security personnel).

Having each moving object report its location to a central server has the potential for generating very frequent updates when the population of objects is large. Thus techniques that accomplish tracking in an efficient manner, without generating unnecessary data communication and without overloading the server with updates, are important.

We employ a shared-prediction-based approach where the server shares a prediction of each moving object's near-future position with the object. The object then monitors the deviation between its predicted position as known by the sever and its actual position (using GPS). An update is sent to the server when needed in order to maintain an agreed-upon accuracy. Thus, the better the predictions, the better the system performs. Therefore, focus is on making robust and accurate near-future position predictions.

The importance of techniques that enable the tracking of moving objects is confirmed by the increasing number of scientific publications on the subject (e.g., [6, 14, 16]) and by the scores of companies [9] that offer commercial tracking services. However, these commercial solutions are almost all time-based (updates are sent to the server at regular time intervals), and no solution known to the authors offers accuracy guarantees at competitive costs.

The future position of an object may be predicted by a *trajectory* [10, 13, 15, 17] that is a polyline in 3-dimensional [13] or 4-dimensional [15] space, the dimensions being a two-dimensional "geographical" space, a time dimension, and, possibly, an uncertainty thresholds dimension.

A trajectory may be computed by using the speed limits and average speeds on the specific road segments on which an object is expected to move. Xu and Wolfson [17] use the average real-time speeds reported every 5 minutes by in-road sensors. Ding and Güting [4] have recently discussed the use of what is essentially segment-based tracking (to be covered in more detail shortly) within an envisioned system based on their own proposal for a data model for the management of road-network constrained moving objects. In our techniques, the prediction of an object's movement is done using the speed and movement direction received from the object. For more advanced prediction, we are able to use acceleration profiles extracted from past movements of the objects.

When only low accuracies of the predicted positions are needed, cellular techniques [1, 11, 12] may be used. In this approach, updates are handled in the cellular mobile network. In contrast to these techniques, we consider much more accurate tracking. We also employ knowledge of the road infrastructure.

Fox et al. [5] explore the use of statistical methods, e.g., multiple hypothesis tracking, in a more abstract location estimation context than the one we consider. Integration of such methods into our setting may enable further analyses.

The TRAX system implements tracking techniques reported on in two previous papers [2, 3]. The demonstrations of the techniques described in this paper illustrates aspects of tracking that could not be presented in these papers. More specifically, the demonstrations show visually how the tracking techniques are influenced by real-world GPS and digital road network inaccuracies, as well as by the actual communication and computation delays. In short, the demonstrations convey a visual understanding of the functioning of the tracking techniques in real-world environments. The demonstrations aim to be of interest to all who do research on the indexing and query and update processing for moving objects.

The remainder of the paper offers an overview of the tracking system, describes two demonstrations, and points to resources.

VLDB '07, September 23–28, 2007, Vienna, Austria.

2. TRACKING SYSTEM OVERVIEW

We assume that the moving objects have computational capabilities, are location aware, and can communicate wirelessly with a central server. In particular, we will assume that a moving object is an individual equipped with a mobile phone that uses GPS for positioning and GPRS or 3G for communication. The system accommodates mobile phones ranging from simple, modern phones to smart phones/pocket PCs. Figure 1 illustrates tracking in this setting using a Pocket PC. Additionally, we assume that different

Figure 1: Tracking of a Pocket PC

objects are to be tracked with different, guaranteed accuracies.

To reduce the communication and update load, each object shares a prediction of its future position (a function from time to space) with the server. When the tracking of an object starts, the object and server agree on the accuracy of the tracking and on the prediction function. As the tracking proceeds, the object continually (each second) compares its GPS position with the position obtained from the prediction function. Only when the distance between these is to exceed the accuracy to be guaranteed, the object sends its most recent geo-location data (*location*, *speed*, *direction*, *time*) to the server. The server then updates its prediction function for the object, and if the object is not able to "guess" this function, the server sends the function to the object. This continues until the tracking is terminated.

The challenge is then how to predict the future position of an object so that the number of updates is kept as low as possible. The reduction in updates reduces the communication cost and the server-side update processing. The system utilizes three prediction techniques.

point-based tracking The prediction is that the object remains at the position reported most recently. This technique works well in certain settings—imagine kids playing soccer being tracked with an accuracy of 200 meters.

vector-based tracking The prediction is that the object moves at constant speed in a fixed direction, according to the most recently reported speed and direction, starting from the most recently reported position at the time that position was observed. So the prediction function is a linear function of time.

segment-based tracking Here, a digital road network is utilized: the prediction is that the object will move at constant speed along the road segment (a polyline) that it is located on. The most recently reported position is map matched on to a segment to obtain the start position for the prediction, and the most recently reported speed is assumed. When the predicted position reaches the end of its road segment, the predicted position remains there (this will eventually trigger an update and thus map matching onto another segment).

Segment-based tracking may utilize pre-recorded routes, which are simply long polylines, and associated acceleration profiles that make it possible to advance an object along a segment in a more sophisticated manner. For each of the three prediction techniques, including segment-based tracking with routes and acceleration profiles, an update needs only consist of (*location*, *speed*, *direction*, *time*). When routes and acceleration profiles are not available, the system simply does not use them. Further, when no digital road network is available or map matching fails, the system switches automatically to vector-based tracking, which is unaffected by such problems.

Figure 2 shows the performance of each tracking technique. The results presented in the figure are obtained from real GPS data from the INFATI project [8]. The dataset used for segment-based tracking with routes and acceleration profiles consists of approximately 57,000 GPS positions, while the dataset used for all the other experiments covered consists of 500,000 GPS positions.

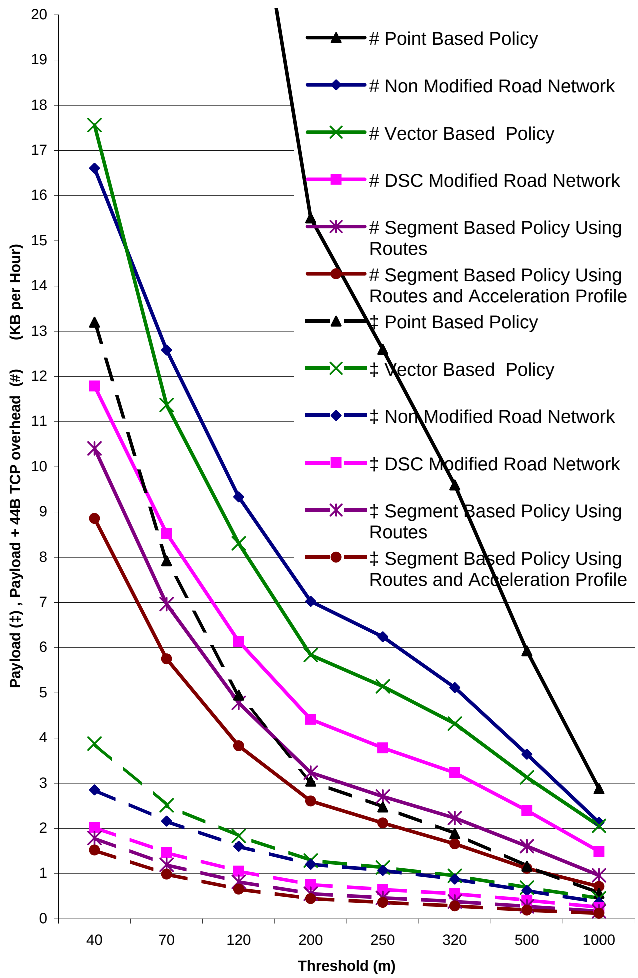

Figure 2: Payload

The x axis presents different accuracy thresholds that vary from 40 to 1000 meters.The y axis presents the average payload in kilobytes per hour of non-stop driving.

Solid curves with labels starting with # report the payload with the TCP overhead (44-byte headers), while the dashed curves starting with ‡ report the payload without any network overhead. Apart from the curves for point-based and vector-based tracking, all curves use segment-based tracking, but with various modifications and improvements.

Point-based tracking is the biggest consumer of bandwidth. The figure does not show payloads bigger than 20 kB. For the accuracies 40, 70, and 120 meters, point-based tracking consumes 67, 40, and 25 kB per hour, respectively.

When using segment-based tracking with a standard digital road network, the performance improves very substantially. In the network used, each part of a road in-between two intersections is represented by a segment. When applying the DSC modification to the road network [3], longer segments are obtained, and segment-based tracking with the resulting road network is now better than vector-based tracking.

Even better performance is achieved using segment-based tracking with routes. Here, routes, which are long polylines, are obtained from the past driving of the vehicles and are used instead of the road-network segments. The best performance is achieved when using routes together with acceleration profiles associated with the routes. Like the routes, the acceleration profiles are obtained from the past driving of the vehicles.

Figure 3 illustrates the architecture of the tracking system. Each

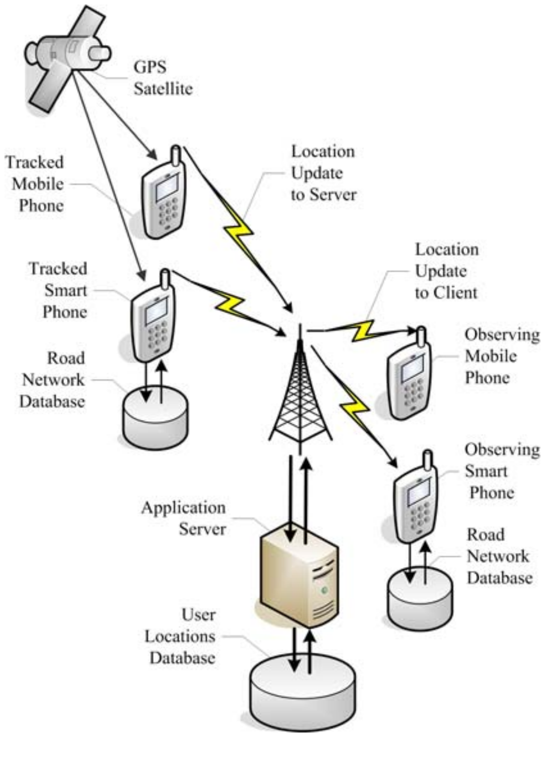

Figure 3: Tracking System Architecture

object being tracked, termed simply a *tracked phone*, is equipped with a mobile phone that has either a built-in GPS receiver or an external GPS receiver, e.g., with Bluetooth connectivity.

Next, each observing object is termed an *observing phone*. Such objects are also equipped with mobile phones that observe, or display, the current locations of one or more tracked phones. They have no need for geo-positioning capabilities.

Finally, the *application server* is aware of which observing phones observe which tracked phones. The server uses this information to distribute updates from tracked phones to the appropriate observing phones. The server is also capable of collecting and storing trajectories.

The system uses pocket PCs with built-in GPS receivers (as the one displayed in Figure 1). The system is implemented using C#, .net, and SQL server mobile edition.

The system accommodates two types of tracked and observing phones: phones that do not store the digital road network locally and more powerful phones that do offer this capability. The former phones support only point- and vector-based tracking, while the latter also support segment-based tracking. If an observing phone does not accommodate the road network, the phones it tracks must use either point- or vector-based tracking. If an observing phone does accommodate the road network, but the phone being tracked does not, the observer must track the less powerful phone using point- or vector-based tracking. Otherwise, any of the three tracking techniques can be used.

3. DEMONSTRATION

We proceed to describe two demonstration setups, an extended-duration setup where object movement is pre-recorded and a more intensive, short-duration setup where video is used to visualize the behavior of the tracking system.

3.1 Extended-Duration Demonstration

The purpose of this demonstration is to illustrate how moving objects can be tracked using the different tracking techniques, as well as to offer detailed insight into the functioning of the tracking techniques in a real-world setting and to demonstrate the performance of each technique.

The demonstration covers point-, vector-, and segment-based tracking, and it makes it possible to observe and compare in real-time the communication and update performance of the three techniques. In addition, the demonstration offers a visualization of each technique that shows how the technique behaves on both the tracked phones and the observing phones.

The demonstration setup, shown in Figure 4, consists of three mobile phones, each running a different tracking technique. In order to compare the techniques, the phones will use the same accuracy threshold and will use identical streams of so-called NMEA sentences with GPS data.

The tracked phones use identical GPS data, but each is utilizing a different tracking technique. They display their estimated positions, as well as the positions obtained from the (simulated) GPS receiver and the maximum allowed inaccuracy.

The server receives updates from the tracked phones and displays the real-time communication and update statistics for each phone. The server stores the received data and sends updates to each of the three observing phones, where each observing phone observes a different tracking phone. An observing phone displays the estimated position of the tracked phone as well as the allowed inaccuracy.

3.2 Video-Enhanced Demonstration

The setup of this demonstration is illustrated in Figure 5. Its goal is to show how the best-performing technique (segment-based tracking) behaves in a real-world environment with GPS and digital road network inaccuracies, as well as communication and computation delays. To achieve this, a person is driving a car equipped

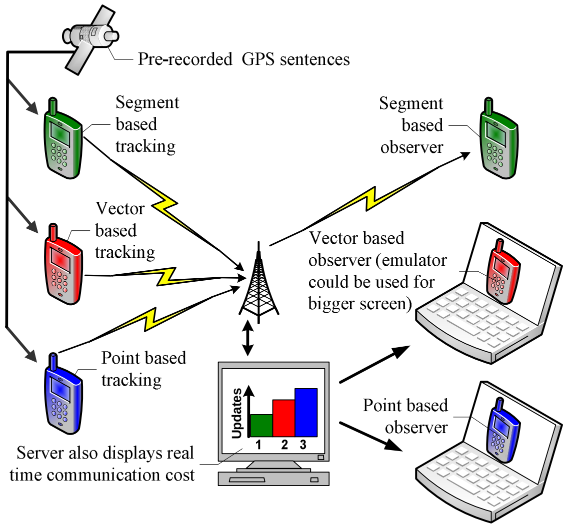

Figure 4: Extended-Duration Demonstration Setup

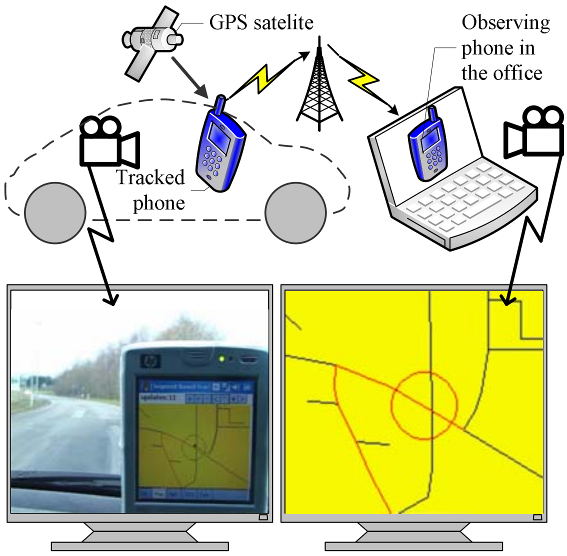

Figure 5: Video-Enhanced Demonstration Setup

with a tracking phone and a camcorder that captures the display of the tracked phone and the driver's view of the road ahead. Simultaneously, another camcorder (screen capturing software) is used to record the screen of the observing phone, which is located in an office.

This demonstration displays the estimated position and the allowed accuracy on the observing phone. The observing phone is being emulated in order to be able to view the phone's display on a large screen; otherwise, a real phone could have been used. The demonstration also encompasses a video recording of the tracked phone, which displays the position obtained from the GPS receiver and the road in front of the car, which shows where the car actually is at the current time. Finally, the demonstration displays the updates sent from the tracked phone to the server and then to observing phone.

4. TO PROBE FURTHER

The tracking techniques covered here are being integrated as a web service into the Streamspin system. This system aims to support easy creation and sharing of mobile services—stated briefly, it aims to provide the "same" functionality as does YouTube, but for mobile services instead of video clips [7]. The Streamspin system is open to the general public, at `streamspin.com`. Thus, anybody may build services that utilize the tracking techniques presented here.

The data used for the experiments covered in the paper are also publicly available [8].

5. REFERENCES

[1] I. F. Akyildiz and J. S. M. Ho. A Mobile User Location Update and Paging Mechanism Under Delay Constraints. *ACM-Baltzer Journal of Wireless Networks*, 1: 244–255 (1995)

[2] A. Čivilis, C. S. Jensen, J. Nenortaite, and S. Pakalnis. Efficient Tracking of Moving Objects with Precision Guarantees. In *Proc. MobiQuitous*, pp. 164–173 (2004)

[3] A. Čivilis, C. S. Jensen, and S. Pakalnis. Techniques for Efficient Road-Network-Based Tracking of Moving Objects. *IEEE TKDE*, 17(5): 698–712 (2005)

[4] Z. Ding and R. H. Güting. Managing Moving Objects on Dynamic Transportation Networks. In *Proc. International Conference on Scientific and Statistical Database Management*, pp. 287–296 (2004)

[5] D. Fox, J. Hightower, L. Liao, D. Schultz, and G. Borriello. Bayesian Filters for Location Estimation. *IEEE Pervasive Computing*, 2(3): 24–33 (2003)

[6] R. H. Güting and M. Schneider. Moving Objects Databases. Morgan Kaufman (2005)

[7] C. S. Jensen. When the Internet Hits the Road. In *Proc. BTW*, pp. 2–16 (2007)

[8] C. S. Jensen, H. Lahrmann, S. Pakalnis, and J. Runge. The INFATI Data. Aalborg University, TimeCenter TR-79. `www.cs.aau.dk/TimeCenter` (2004)

[9] C. S. Jensen, K. -J. Lee, S. Pakalnis, S. Šaltenis. Advanced Tracking of Vehicles. In *Proc. European Congress and Exhibition on ITS*, 12 pages (2005)

[10] H. A. Karimi and X. Liu. A Predictive Location Model for Location-Based Services. In *Proc. ACM International Symposium on Advances in Geographic Information Systems*, pp. 126–133 (2003)

[11] G. Li, K. Lam, and T. Kuo. Location Update Generation in Cellular Mobile Computing Systems. In *Proc. International Parallel & Distributed Processing Symposium*, p. 96 (2001)

[12] Z. Naor and H. Levy. Minimizing the Wireless Cost of Tracking Mobile Users: An Adaptive Threshold Scheme. In *Proc. IEEE INFOCOM*, pp. 720–727 (1998)

[13] G. Trajcevski, O. Wolfson, F. Zhang, and S. Chamberlain. The Geometry of Uncertainty in Moving Objects Databases. In *Proc. EDBT*, pp. 233–250 (2002)

[14] A. Voisard, J. Schiller. Location-Based Services. Morgan Kaufmann (2004)

[15] O. Wolfson. The Opportunities and Challenges of Location Information Management. In *Proc. Intersections of Geospatial Information and Information Technology Workshop* (2001).

[16] O. Wolfson and H. Yin. Accuracy and Resource Consumption in Tracking and Location Prediction. In *Proc. SSTD*, pp. 325–343 (2003)

[17] B. Xu and O. Wolfson. Time-Series Prediction with Applications to Traffic and Moving Objects Databases. In *Proc. ACM MobiDE*, pp. 56–60 (2003)