ADVANCED TRACKING OF VEHICLES

Christian S. Jensen[†] Ki-Joune Li[‡] Stardas Pakalnis[†] Simonas Šaltenis[†]

[†]Department of Computer Science, Aalborg University, Denmark
[‡]Department of Computer Science, Pusan National University, South Korea

ABSTRACT

With the continued advances in wireless communications, geo-location technologies, and consumer electronics, it is becoming possible to accurately track the time-varying location of each vehicle in a population of vehicles. This paper reports on ongoing research that has as it objective to develop efficient tracking techniques. More specifically, while almost all commercially available tracking solutions simply offer time-based sampling of positions, this paper's techniques aim to offer a guaranteed tracking accuracy for each vehicle at the lowest possible costs, in terms of network traffic and server-side updates. This is achieved by designing, prototyping, and testing novel tracking techniques that exploit knowledge of the road network and past movement. These resulting tracking techniques are to support mobile services that rely on the existence of a central server that continuously tracks the current positions of vehicles.

INTRODUCTION

The continued improvements in wireless communications and geo-positioning as well as improvements in the performance/price ratio for consumer electronics combine to enable a new infrastructure that enables a kind of online mobile service that relies on the tracking of the continuously changing positions of entire populations of service users. This type of service is characterized by large volumes of updates, which in turn gives prominence to the development of techniques for advanced location prediction and update reduction.

A range of applications may utilize the tracking of the positions of vehicles. Such applications include mobile services in relation to traffic monitoring, collective transport, and the management of fleets, e.g., of emergency vehicles, police cars, delivery trucks, and vehicles carrying dangerous or highly valuable cargo. For example, traffic jams may be identified by monitoring the movements of vehicles; and the service users that should receive specific traffic-jam or other traffic-related information are identified by tracking the users' positions. Some services depend only on fairly inaccurate tracking, e.g., weather information services, while other services require more accurate tracking.

The present paper reports on research that addresses the following problem: given a population of vehicles, where the position of each vehicle must be known at some specified accuracy, provide techniques that enable a central server to continuously and efficiently maintain the position of each vehicle within the vehicle's specified accuracy.

With few exceptions, current commercial tracking offerings simply report a position sample to the central server every n time units for each vehicle. Some products are also capable of reporting a vehicle's position when the direction of the vehicle changes. These approaches do not offer accuracy guarantees.

The desired tracking techniques are to be robust in the sense that they must not depend for their functioning on the availability of map data and successful map matching. The desired techniques must be efficient in the sense that they reduce system load, including network communication and server update loads.

To the best of the authors' knowledge, this paper reports on the currently best such techniques. The objective has been achieved by first developing and testing a variety of tracking techniques. These techniques are then combined into a single, robust tracking technique. This technique is then implemented as a tracking component. The resulting component may be used by a variety of mobile services and applications, as mentioned above.

The presentation is structured as follows In the next section, we proceed to describe the overall tracking framework and the specific tracking techniques. This is followed by a section, "Implementation and Evaluation," that offers descriptions of two implementations of the techniques—one that is used for efficiency evaluation and one that demonstrates proof of concept. Then follows a section that covers related developments in some detail, in both the commercial and research arenas. A final section summarizes the paper.

TRACKING FRAMEWORK AND TECHNIQUES

We first describe the general tracing framework. We then give an overview of three basic tracking techniques and cover improvements to one of these.

Tracking Framework

We assume that the vehicles are online via some form of wireless communication network. In particular, we have utilized GPRS connections. We also assume that the positions of the vehicles are available. Specifically, we rely on the Global Positioning System (GPS) for positioning. Thus vehicles are simply assumed to be equipped with the equivalent of a modern mobile phone and a GPS receiver. We also assume that the vehicles travel in a road network.

To accomplish tracking with a certain accuracy, each vehicle is aware of where the server thinks it is currently located. Each vehicle repeatedly samples its real (GPS) position and compares this with the position assumed by the sever. When needed in order to maintain the required accuracy, the vehicle issues an update to the server. The server may predict the future positions of a vehicle in different ways—we shortly elaborate on this. Having received an update from a vehicle, the server in turn explicitly informs the vehicle about how it now predicts the vehicle's position. The challenge is then how to represent, and predict, the future positions of a vehicle so that the number of updates is minimized. Reduction of updates reduces communication and server-side update processing.

Figure 1 presents a diagram describing the tracking framework. The client initially obtains its location information from the GPS receiver. It then establishes a connection with the server and issues an update, sending its GPS information and unique identity to the server.

Having received this update, the server determines which tracking technique and threshold to use for the client (these are predefined), and it stores the information received from the client in its database. We assume that the tracking threshold exceed the typical accuracy offered by the GPS receiver. In the experiments we report on later in the paper, the lowest threshold is 40 meters. If the tracking technique is the segment-based one, the server also determines on which road segment the client is moving. The server then sends its representation of the client's position to the client.

Having received and stored this information from the server, the client obtains its current location information from the GPS receiver. The client then calculates its predicted position and compares this to the GPS position. If the difference between these two exceeds the specified threshold, the client issues an update to the server. If not, a new comparison is made. This procedure continues until it is terminated by the client.

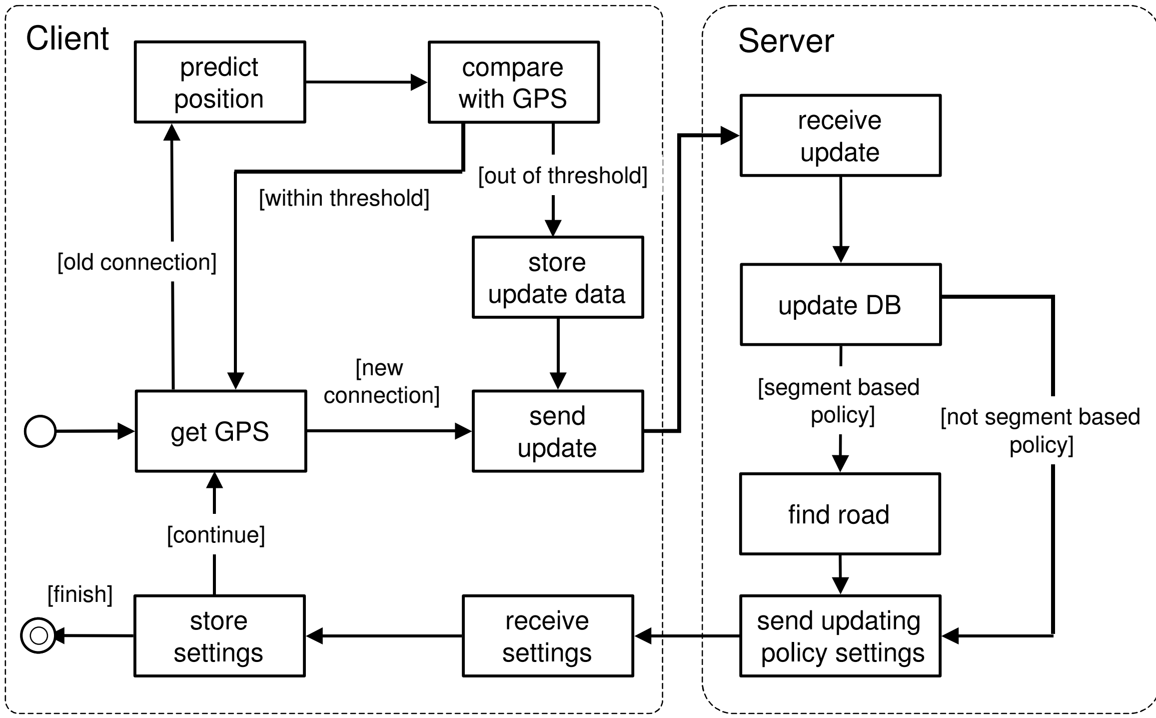

Figure 1: Tracking Framework

Point-, Vector-, and Segment-Based Tracking

We proceed to present different kinds of tracking techniques that may be used withing the tracking framework.

With *point-based tracking*, the server predicts that a vehicle's future positions are identical to the most recently reported position. An update is issued by a vehicle when its distance to the previously reported position deviates from its current GPS position by the specified accuracy. Figure 2(a) exemplifies point-based tracking. Here, points indicate updates. The solid line represents the road, and the thin line represents the actual movement of the vehicle.

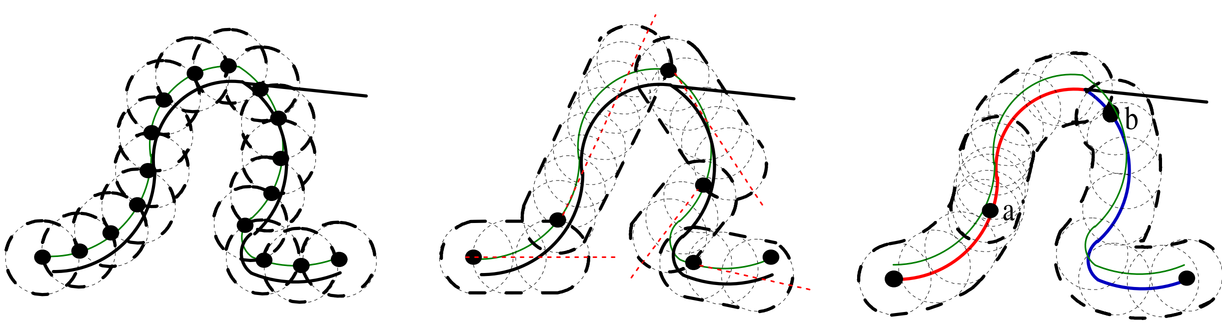

a) Point-Based b) Vector-Based c) Segment-Based

Figure 2: Tracking Policies

Next, with *vector-based tracking*, the future positions of a vehicle are given by a linear function of time, i.e., by a start position and a velocity vector. The specific start position and velocity vector used are those given by the GPS receiver in the update most recently received from the vehicle. Point-based tracking corresponds to the special case where the velocity vector is the zero-vector. GPS receivers compute both the speed and heading of the vehicles they are associated with—the velocity vector used in this representation is computed from these two. See Figure 2(b).

Third, in *segment-based tracking*, the main idea is to utilize knowledge of the road network in which the vehicles are moving. A digital representation of the road network is required to be available. The server uses the location information it receives from a vehicle to locate the road segment, represented as a polyline, on which the vehicle is moving. In segment-based tracking, the future positions of a vehicle are given by movement at constant speed along the identified segment. The speed used is the speed most recently reported by the vehicle. If the predicted position reaches the end of a segment, the predicted position remains at the end from then on.

Segment-based tracking is sensitive to the fidelity of the road network representation used. If, for some reason, a matching road segment cannot be found when a vehicle issues an update, the segment-based approach switches temporarily to the vector-based approach, which is always applicable. On the next update, the server will again try to find a matching road segment in the database. Note also that the policy effectively switches to point-based tracking when a vehicle reaches the end of a segment.

An example of segment-based tracking is shown in Figure 2(c). Notice that update points are now located on the road segment, not on the trajectory obtained via the GPS receiver.

Additional technical details on the three tracking techniques covered above are available elsewhere [3, 4].

Improved Segment-Based Tracking

Segment-based tracking is the most promising of the three tracking techniques. To be robust, it fundamentally exploits vector-based tracking, and it also uses point-based tracking at the ends of road segments. Further, it positions vehicles with respect to the transportation infrastructure, which is important for many mobile services because content used by these services is also positioned with respect to the infrastructure. A variety of variants of segment-based tracking have been developed and tested.

The first variant modifies the road network representation so that vehicles stay longer on the same segments. In particular, the modifications give priority to merging smaller road segments that belong to the same named street or road, they attempt to obtain straight segments, and they try to avoid dead-end or near dead-end segments when merging. Creation of longer road segments in the road network representation is important because an update intuitively occurs each time a vehicle reaches the end of a segment.

The second variant anticipates the current route followed by a vehicle. Briefly, a route is nothing but a long polyline obtained by connecting polylines that represent (possibly partial) segments of the road network. The use of routes offers substantial reduction in the numbers of updates caused by segment changes, but, of course, relies on the availability of routes. Routes may be obtained from past movement [2], if the vehicle is driving in familiar surroundings, or by a navigation system, which may be in use if the vehicle is driving in unknown surroundings.

The third variant attaches acceleration profiles to routes. The previous approaches assume that vehicles move at constant speed in-between updates. In order to reduce the number of updates caused by the speed variations of a vehicle, more precise speed modeling is introduced.

Additional technical details on the improvements described here are available in the literature [5, 6].

IMPLEMENTATION AND EVALUATION

Two implementations of the tracking techniques were carried out, each with its own purpose. The first is an implementation based on Oracle/PLSQL. Its purpose is to enable effective performance evaluation of the techniques using already collected logs of GPS data from vehicles. The second implementation was conducted on a mobile-phone platform with the purpose of gaining experience with tracking in a representative, actual infrastructure within which the tracking is to be applied. We describe each in turn.

Performance Evaluation of Tracking

Using an Oracle/PLSQL implementation, the six tracking techniques covered in the previous section were evaluated. Specifically, a road network representation and GPS logs from vehicles that participated in a speed adaptation project were used for simulation [10]. Approximately 458,000 GPS positions from five cars were used, and accuracies ranging from 40 to 1000 meters were investigated.

Figure 3 illustrates the uses of data in the experiments carried out for each technique using the Oracle/PLSQL implementation. To evaluate the point-based and vector-based tracking techniques, only

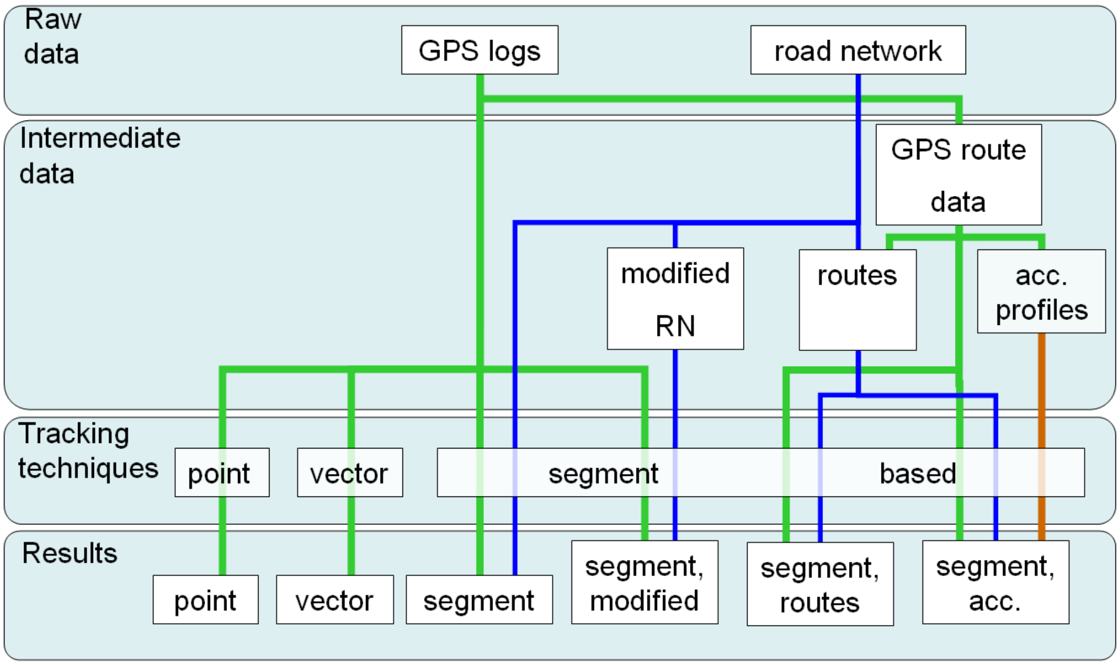

Figure 3: Data Usage in Experimental Performance Study

input from the GPS logs was used. The experiments with segment-based tracking with the unmodified road network used the raw road network representation as and additional input.

The experiments with the improvements to segment-based tracking used additional data obtained from the GPS logs and the raw road network. The extra input for the experiment with segment-based tracing using a modified road network was obtained by connecting road segments of the raw road network, as described in the previous section. For the experiments using routes and the experiments using routes with acceleration profiles, routes from home to work of one driver were used. Here, we first identified trips from home to work ("GPS route data"), and then obtained routes by connecting (partial) road segments that were covered by GPS route data. The GPS route data consists of 56,000 GPS log records. For segment-based tracking using routes with acceleration profiles, we have additionally extracted acceleration profiles from the GPS route data.

Representative results, presented in Figure 4, show that depending on the tracking technique and threshold, the average time in-between updates from a vehicle varies from 3 to 280 seconds. Point-based tracking has the worst performance for all thresholds. Vector-based tracking is better than basic segment-based tracking. This is because segments are short in the unmodified road network. Vector-based tracking is inferior to any of the improved variants of segment-based tracking. More specifically, segment-based tracking using a modified road network outperforms vector-based tracking. The use of

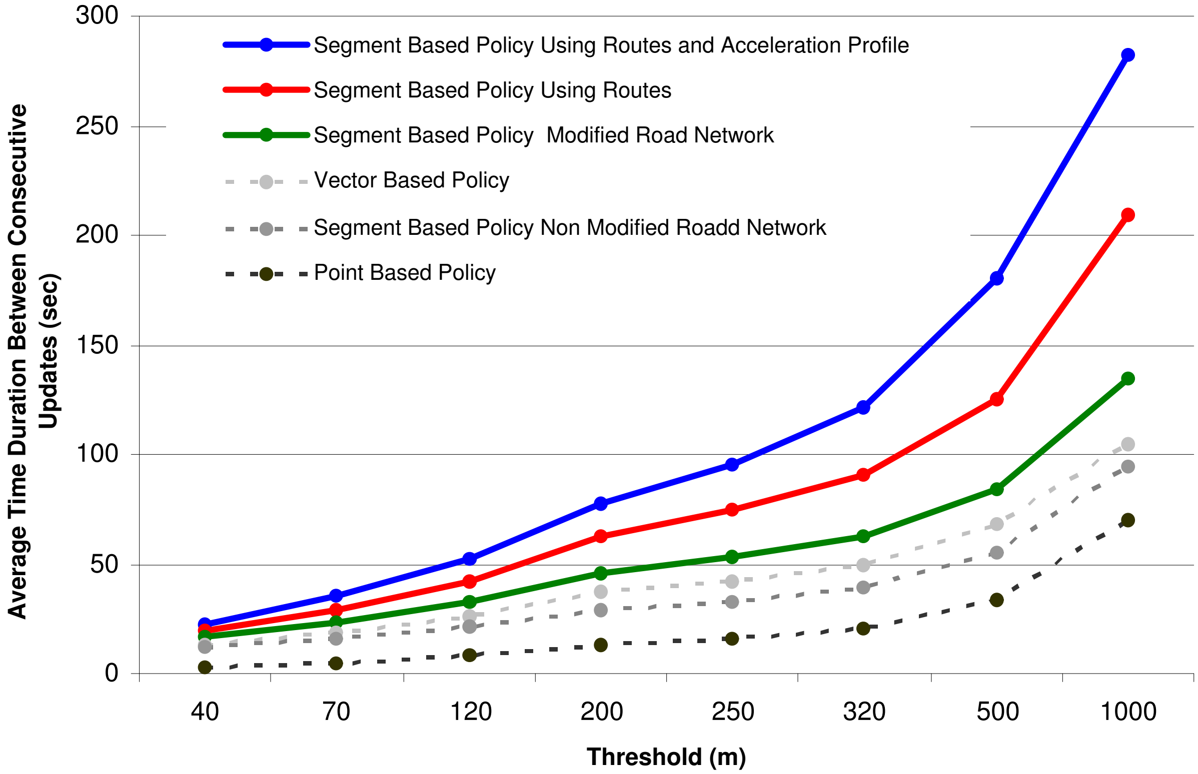

Figure 4: Performance Comparison of Tracking Techniques

routes leads to further improvements. Segment-based tracking using routes with acceleration profiles offers the best results. We believe that the update rates reported here demonstrate that it is feasible to perform quite accurate tracking in fleet-management applications.

Proof of Concept Implementation

To test the tracking framework and techniques in the actual infrastructure in which they are intended to be used, we implemented the tracking techniques using a typical, modern mobile phone with a standard GPRS subscription, a standard GPS receiver, and standard server-side hardware and software.

Figure 5 shows the client-side hardware on which the tracking techniques were implemented. The implementation will work with any programmable mobile phone with an external or internal GPS receiver.

The architecture of the system is illustrated in Figure 6. Here, the mobile client receives location data (specifically, NMEA sentences) from the GPS receiver, and it communicates with the server through the GSM network using the http protocol.

The client employs a scalable vector graphics viewer together with its tracking service. We have used a framework [9] that other location-based services can be plugged into. The use of a single framework for all location-based services enables better use of shared resources (e.g., the GPS receiver or cached maps). In case of users that use only the tracking service, tracking component can be implemented as a separate application instead of as a plug-in service within the framework.

The architecture enables the use of a thin client that is unable to store the entire road network. Thus the main logic is placed on the server side. Operations like map matching are carried out in Oracle using stored procedures, thus reducing the communication between application server and the database management system. If the user needs maps for use with the tracking or other services, these can be obtained from an external web map server and then sent to the client together with other data.

Figure 5: Tracking Hardware

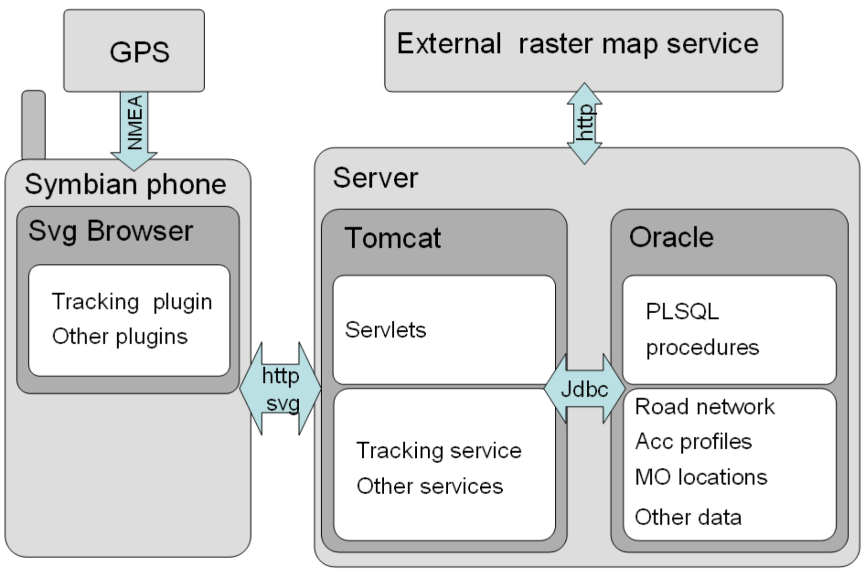

Figure 6: Tracking Application Architecture

The computing power of the Nokia 7650 used in this implementation was sufficient for the tracking. Specifically, the tracking was able to run in the background, allowing the user to perform other tasks, including displaying an animation of the tracking.

Since the processing power and memory capacities of mobile phones continue to increase, it is becoming relevant to also take into consideration scenarios where the phone is able to store the entire road network or cache significant part of it. The tracking framework with segment-based tracking will benefit from the use of a phone that stores the road network, since sending identities of road segments, rather than also their geometries, is sufficient. In addition, with more processing power, the phone may perform the map matching instead of the server, thus distributing computation and reducing the server load.

RELATED DEVELOPMENTS

We proceed to offer an overview of related developments in the commercial and academic arenas. We first offer an overview of 21 tracking-related products and services that we believe are representative of the current commercial state of the art. We then provide an overview of related works within the academic community that may impact future commercial offerings, and it briefly points to the differences between these and the techniques discussed earlier in this paper.

Commercially Available Products and Services

Table 1 summarizes pertinent properties of what we believe is a representative range of commercially available tracking solutions. The table captures properties of 21 solutions provided by 19 companies. The information that went into the creation of the table has been obtained via the Internet during February and March, 2005.

Company	**Product**	GPS	Cell	WAAS	SMS	GPRS,CDPD	Satellite	Phone	Custom	Time based	On request	Sensors	Spatial events
BSM	Sentinel	√		√		√			√	√		√	√
Cellfind	look4me		√		√			√			√		
Cybit	mapAmobile		√		√			√			√		
Euman	LifePilot	√				√		√		√			
Fleetella	FL1700	√				√			√	√	√		
FleetOnline	FleetOnline		√		√			√			√		
	Trimtrac	√			√				√	√	√		
Global Tracking Solutions	GTS-1000	√				√			√	√		√	√
	Sat-TDiS	√					√		√	√		√	√
GPS Fleet Solutions	Marcus	√				√			√	√		√	√
Gpsnext	Stealth tracker	√				√			√	√			
Guard Magic	VS, VG	√				√	√		√		√	√	
Mapbyte	Mapaphone		√		√			√			√		
Mobitrac	Mobitrac	√				√	√		√	√			√
Siemens	m.traction Senior Care Service		√		√				√		√		
Telus	Action Tracker	√				√		√		√			
uLocate	fleeTracker	√				√		√		√			
Unteh	Mobitrack	√				√		√		√			
Veriloaction	VL-Tracer	√				√			√	√			
Web Tech wireless	WebTech5000	√				√			√	√		√	√
2020 Fleet Management	Sentinel Live	√				√			√	√	√		

Table 1: Properties of Tracking Solutions

The first column lists company names, and the second lists product names. Starting from the third, each column concerns one product property, and a check mark in a cell indicates that the product in the row of the cell possesses the property corresponding to the column of the cell. No check mark indicates the opposite.

Columns *GPS*, *Cell* and *WAAS* concern the means of positioning supported, with *Cell* denoting cellular network based positioning and *WAAS* denotes the wide-area augmentation system which is based on GPS, but offers higher accuracy than GPS by using corrections. The next three columns concern the types of communication supported, with *SMS* denoting the short messaging service, *GPRS/CDPD* denoting general packet radio service (GSM based) and cellular digital packet data, and *Satellite* denoting satellite-based communications. Then two columns follow that capture the types of terminals supported, with *Phone* denoting mobile phones and *Custom* denoting custom terminals. Finally, the last four columns characterize the type of tracking or how position updates are generated. Here, *Time based* denotes time-based tracking, i.e., updates are issued at regular time intervals, *On request* means that positions of moving objects are pulled from the clients only on request, *Sensors*, means that position updates can be generated according to input from external sensors, e.g., an alarm, a speedometer reading, a thermometer reading, and *Spatial events* means that position update can be generated by the object entering or leaving a certain region, e.g., when leaving the city limits or a pre-specified route or when getting within a certain range of a point of interest.

It should be noted that none of the products described in Table 1 provides accuracy guarantees or support accuracy based tracking. Advanced options such as *Spatial events* are usually supported by solutions involving large, custom-made terminals.

Related Research

When predicting the future position of an object, the notion of a *trajectory* is typically used [11, 15, 16, 19], where a trajectory is defined in 3-dimensional [15] or 4-dimensional [16] space. The dimensions are a two-dimensional "geographical" space, a time dimension, and an uncertainty thresholds dimension. A point in this space then indicates, for a point in time, the location of an object and the uncertainty of the location. Such points may be computed using speed limits and average speeds on specific road segments belonging to a trajectory. Xu and Wolfson [19] use average real-time speeds reported every 5 minutes by in-road sensors. In contrast, the techniques presented in this paper use the actual speed received from an object for predicting the object's movement. For more accurate prediction, we use acceleration profiles that quite accurately model speed variations along routes. Such profiles are influenced by both the properties of the actual roads being traversed, the degree of congestion, and the habits of a concrete driver.

Wolfson et al. [17] propose tracking techniques that offer accuracy guarantees. These assume that objects move on predefined routes already known to the objects, and route selection is done on the client side. If an object changes its route, it sends a position update with information about the new route to the server. The techniques described in this paper improve on this by also accommodating objects with memory restrictions, and they also work in cases where routes are not known or where map matching does not succeed.

Lam et al. [12] present an adaptive monitoring method that takes into consideration the update, deviation, and uncertainty costs associated with tracking. The method also takes into account the cost of providing incorrect results to queries, during the process of determining when to issue updates. With this method, the moving objects that fall into a query region need close monitoring, and a small accuracy threshold is used for them. Objects not inside a query region may have big thresholds. The techniques presented in this paper allow different objects to have different thresholds, and they allow thresholds to change dynamically.

A proposal for trajectory prediction by Karimi and Liu [11] assigns probabilities to the roads emanating from an intersection according to how likely it is that an object entering the intersection will proceed on them. The sub-road network within a circular area around an object is extracted, and the most probable route within this network is used for prediction. When the object leaves the current sub-network, a new sub-network is extracted, and the procedure is repeated. The probabilities are not individual to each object. Rather, the same probabilities are used for all objects, and they do not take into account past choices during the trip of an object. This contrasts our use of routes and acceleration profiles.

Wolfson et al. [19] have recently investigated how to incorporate travel-speed prediction in a database. They assume that sensors that can send up-to-date speed information are installed in the roads. In contrast, we use so-called GPS-based floating-car data, and we predict positions based on historical records and for each moving object in isolation. This avoids the need for in-road sensors and for gathering information from such sensors.

Next, Wolfson and Yin [18] consider tracking with accuracy guarantees. Based on experiments with synthetic data, generated to resemble real movement data, they conclude that a version of the point-based tracking, as discussed in the earlier section covering point-, vector, and segment-based tracking, is outperformed by a tracking technique that resembles the segment-based tracking also discussed in that earlier section. For a small threshold of 80 m, the latter is a bit more than twice as good as the former; for larger thresholds, the difference decreases. Their dependent variable is numbers of updates per distance unit. They consider neither the use of routes nor acceleration profiles and assume that map matching is always successful.

It should also be noted that Ding and Güting [7] have recently discussed the use of what is essentially segment-based tracking within an envisioned system based on their own proposal for a data model for the management of road-network constrained moving objects.

Gowrisankar and Nittel [8] introduce briefly a dead-reckoning policy that uses angular and road deviations, so that an update is issued whenever one of these deviations exceeds a defined threshold.

When only low accuracy of predicted positions are needed, cellular techniques [1, 13, 14] may be used. With such techniques, the mobile network tracks the cells of the mobile objects in real time in order to be able to deliver messages or calls to the objects. In this approach, update is handled in the mobile network. In contrast to these techniques, we assume scenarios where higher accuracies, well beyond those given by the cells associated with the base stations in a cellular network, are needed and where positioning with respect to a road network is attractive.

SUMMARY

This paper offers an overview of a framework for tracking of vehicles with accuracy guarantees. It describes three basic tracking techniques that work within this framework, namely point-, vector-, and segment-based tracking; and it describes three improvements to segment-based tracking. One implementation, with the purpose of evaluating the performance of the tracking framework and techniques, is reported; and another, developed with the objective of demonstrating that the kind of tracking proposed here is feasible using standard technologies is also reported. The paper describes results obtained via the first implementation, and it describes the architecture of the second implementation. Finally, the paper offers an overview of commercially available tracking-related solutions, and it describes related work from the research community.

ACKNOWLEDGMENTS

This work was supported in part by grants from Center for ITS, Aalborg University, Denmark; the Nykredit Corporation; the Electronics and Telecommunications Research Institute, South Korea; and the Danish National Center for IT Research.

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. Brilingaitė, C. S. Jensen, and N. Zokaitė. Enabling Routes as Context in Mobile Services *ACM International Workshop on Geographic Information Systems*, pp. 127-136, 2004.

[3] A. Čivilis, C. S. Jensen, J. Nenortaite, and S. Pakalnis. Efficient Tracking of Moving Objects with Precision Guarantees. *Proceedings of the First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services*, pp. 164–173, 2004.

[4] A. Čivilis, C. S. Jensen, J. Nenortaitė, and S. Pakalnis. Efficient Tracking of Moving Objects with Precision Guarantees. DB Technical Report TR-5, Department of Computer Science, Aalborg University, 2004, 23 pages.

[5] A. Čivilis, C. S. Jensen, and S. Pakalnis. Techniques for Efficient Tracking of Road-Network-Based Moving Objects. *IEEE Transactions on Knowledge and Data Engineering*, 17(5): 698–712, May 2005 (tentative).

[6] A. Čivilis, C. S. Jensen, and S. Pakalnis. Techniques for Efficient Tracking of Road-Network-Based Moving Objects. DB Technical Report TR-10, Department of Computer Science, Aalborg University, 2004, 37 pages.

[7] Z. Ding and R. H. Güting. Managing Moving Objects on Dynamic Transportation Networks. *proceedings of the International Conference on Scientific and Statistical Database Management*, pp. 287–296, 2004.

[8] H. Gowrisankar and S. Nittel. Reducing Uncertainty In Location Prediction Of Moving Objects In Road Networks. *Proceedings of the Conference on Geographic Information Science*, 2002. <http://www.spatial.maine.edu/~nittel/publications/giscience02_hari.pdf>

[9] F. Hansen, N. O. Bouvin, B. G. Christensen, K. Grønbæk, T. B. Pedersen, and J. Gagach. Integrating the Web and the World: Contextual Trails on the Move. *Proceedings of the ACM Conference on Hypertext and Hypermedia*, pp. 98–107, 2004.

[10] C. S. Jensen, H. Lahrmann, S. Pakalnis, and J. Runge. The INFATI Data. Aalborg University, TimeCenter TR-79, 2004. <http://www.cs.aau.dk/TimeCenter>. Also CoRR cs.DB/0410001 (2004), <http://arxiv.org/abs/cs.DB/0410001>.

[11] H. A. Karimi and X. Liu. A Predictive Location Model for Location-Based Services. *Proceedings of the ACM International Symposium on Advances in Geographic Information Systems*, pp. 126–133, 2003.

[12] K. Y. Lam, O. Ulusoy, T. S. H. Lee, E. Chan, and G. Li. An Efficient Method for Generating Location Updates for Processing of Location-Dependent Continuous Queries. *Proceedings of the International Conference on Database Systems for Advanced Applications*, pp. 218–225, 2001.

[13] G. Li, K. Lam, and T. Kuo. Location Update Generation in Cellular Mobile Computing Systems. *Proceedings of the Workshop on Parallel and Distributed Real-Time Systems*, p. 96, 2001.

[14] Z. Naor and H. Levy. Minimizing the Wireless Cost of Tracking Mobile Users: An Adaptive Threshold Scheme. *Proceedings of the IEEE INFOCOM Conference*, pp. 720–727, 1998.

[15] G. Trajcevski, O. Wolfson, F. Zhang, and S. Chamberlain. The Geometry of Uncertainty in Moving Objects Databases. *Proceedings of the International Conference on Extending Database Technology*, pp. 233–250, 2002.

[16] O. Wolfson. The Opportunities and Challenges of Location Information Management. *Proceedings of the Workshop in Intersections of Geospatial Information and Information Technology*, 2001.

[17] O. Wolfson, A. P. Sistla, S. Camberlain, and Y. Yesha. Updating and Querying Databases that Track Mobile Units. *Distributed and Parallel Databases*, 7(3): 257–287, 1999.

[18] O. Wolfson and H. Yin. Accuracy and Resource Concumption in Tracking and Location Prediction. *Proceedings of the 8th International Symposium on Advances in Spatial and Temporal Databases*, pp. 325–343, 2003.

[19] B. Xu and O. Wolfson. Time-Series Prediction with Applications to Traffic and Moving Objects Databases. *Proceedings of the Third ACM International Workshop on Data Engineering for Wireless and Mobile Access*, pp. 56–60, 2003.