

Techniques for Efficient Road-Network-Based Tracking of Moving Objects

Alminas Čivilis, Christian S. Jensen, and Stardas Pakalnis

March 4, 2005

TR-10

A DB Technical Report

Title Techniques for Efficient Road-Network-Based Tracking of Moving Objects

Copyright © 2005 Alminas Čivilis, Christian S. Jensen, and Stardas Pakalnis. All rights reserved.

Author(s) Alminas Čivilis, Christian S. Jensen, and Stardas Pakalnis

Publication History Extended version of: A. Čivilis, C. S. Jensen, and S. Pakalnis, “Techniques for Efficient Tracking of Road-Network-Based Moving Objects,” to appear in *IEEE Transactions on Knowledge and Data Engineering*, 15 pages.

For additional information, see the DB TECH REPORTS homepage: www.cs.aau.dk/DBTR.

Any software made available via DB TECH REPORTS is provided “as is” and without any express or implied warranties, including, without limitation, the implied warranty of merchantability and fitness for a particular purpose.

The DB TECH REPORTS icon is made from two letters in an early version of the Rune alphabet, which was used by the Vikings, among others. Runes have angular shapes and lack horizontal lines because the primary storage medium was wood, although they may also be found on jewelry, tools, and weapons. Runes were perceived as having magic, hidden powers. The first letter in the logo is “Dagaz,” the rune for day or daylight and the phonetic equivalent of “d.” Its meanings include happiness, activity, and satisfaction. The second letter is “Berkano,” which is associated with the birch tree. Its divinatory meanings include health, new beginnings, growth, plenty, and clearance. It is associated with Idun, goddess of Spring, and with fertility. It is the phonetic equivalent of “b.”

Abstract

With the continued advances in wireless communications, geo-positioning, and consumer electronics, an infrastructure is emerging that enables location-based services that rely on the tracking of the continuously changing positions of entire populations of service users, termed moving objects. This scenario is characterized by large volumes of updates, for which reason location update technologies become important. A setting is assumed in which a central database stores a representation of each moving object's current position. This position is to be maintained so that it deviates from the user's real position by at most a given threshold. To do so, each moving object stores locally the central representation of its position. Then an object updates the database whenever the deviation between its actual position (as obtained from a GPS device) and the database position exceeds the threshold. The main issue considered is how to represent the location of a moving object in a database so that tracking can be done with as few updates as possible. The paper proposes to use the road network within which the objects are assumed to move for predicting their future positions. The paper presents algorithms that modify an initial road-network representation, so that it works better as a basis for predicting an object's position; it proposes to use known movement patterns of the object, in the form of routes; and it proposes to use acceleration profiles together with the routes. Using real GPS-data and a corresponding real road network, the paper offers empirical evaluations and comparisons that include three existing approaches and all the proposed approaches.

Keywords: Database management, distributed databases, query processing, temporal databases

1 Introduction

In step with the emergence of an infrastructure for mobile, on-line location-based services (LBSs) for general consumers, such services are attracting increasing attention in industry and academia.

An LBS is a service that provides location-based information to mobile users. The main idea is to provide the service user with a service that is dependent on positional information associated with the user, most importantly the user's current location. The service may also be dependent on other factors, such as personal preferences and interests of the user [3].

Examples of LBSs abound. A service might inform its users about traffic jams and weather situations that are expected to be of relevance to each user. A friend monitor may inform each user about the current whereabouts of friends. Other services may track the positions of emergency vehicles, police cars, security personnel, hazardous materials, or public transport. A more advanced location-based "catch the monster" game may allow a group of users to work together to surround and catch a virtual, but geo-positioned, monster.

Services such as these rely to varying degrees on the tracking of the geographical positions of moving objects. For example, traffic jams may be identified by monitoring the movements of service users; and the users that should receive specific traffic-jam or weather information are identified by tracking the users' positions. Some services require only fairly inaccurate tracking, e.g., the weather service, while other services require much more accurate tracking, e.g., location-based games.

We assume that users have wireless devices (e.g., mobile phones) that are online via some form of wireless communication network. We also assume that the positions of the users are available. Specifically, we rely on the Global Positioning System for positioning. To accomplish tracking with a certain accuracy, each wireless device monitors its real position (its GPS position) and compares this with a local copy of the position that the central database assumes. When needed in order to maintain the required accuracy in the database, the wireless device issues an update to the server. The database may predict the future positions of a device in different ways. In the general case, the database explicitly informs the mobile device about how it predicts the client's position. The challenge is then how to represent, and predict, the future positions of a mobile device in the database so that the number of updates is minimized. Reduction of updates reduces communication and server-side update processing.

A detailed coverage of related work is given in Section 6. In short, to the best of our knowledge the techniques for update reduction proposed in this paper have not been proposed or evaluated in past work. We share the general setting with Wolfson et al. [16, 18], and our proposals take the segment-based technique described by Čivilis et al. [4], which is similar to a technique presented by Wolfson and Yin [18], as the starting point.

Section 2 describes the segment-based approach in some detail. In this approach, the future movement of a mobile device, termed a moving object, is represented by a road segment drawn from the underlying road network and a fixed speed. A road segment is a polyline, i.e., a sequence of connected line segments. So, this representation assumes that a moving object moves on a known road segment with constant speed.

As explained above, a moving object is aware of the server-side representation of its movement. The server uses the presentation for predicting the current position of the moving object. The client-side moving object uses the representation for ensuring that the server’s predicted position is within the predefined accuracy.

This paper presents techniques that aim to improve the basic segment-based approach. We are basing our proposals on the segment-based approach because we find this to be the most promising outset for more advanced tracking techniques. As an added benefit, by relating moving objects to the underlying road network, we gain easy access to content that is connected to the road network. Such content may be useful in many LBSs. The paper presents the following techniques that improve the segment-based approach:

- Modification of the road network. The number of updates turns out to be closely related to the segmentation of the road network. We present techniques for modification of a road network with the purpose of finding appropriate segmentations of a road network.
- Use of anticipated routes for the moving objects. By using routes in place of segments allows us to reduce the number of updates caused by changes of segments. Routes are represented as (long) polylines.
- Introduction of acceleration profiles. The basic approach assumes that moving objects are moving at constant speed in-between updates. In order to reduce the number of updates caused by the speed variations of the moving objects, we introduce more accurate speed modeling.

In summary, the paper’s main contributions are (i) proposals for three types of techniques that aim to reduce the communication and update costs associated with the tracking of moving objects with accuracy guarantees and (ii) empirical evaluations of the best existing tracking techniques and the new techniques based on real data.

The paper proceeds as follows. Section 2 summarizes previously proposed update policies and offers motivation for further investigation of the segment-based approach. Section 3 covers improvements of segment-based approach using road-network modifications. Sections 4 and 5 present the techniques for update reduction using routes and acceleration profiles, respectively. Section 6 offers an overview of related work. Section 7 summarizes, provides concluding remarks, and offers suggestions for future research. Finally, an appendix reports on additional, detailed experiments with tracking techniques presented in the paper.

2 Background

In this section, we first describe the general tracking scenario that we will use. Then follows a description of the position data used for tracking. Then we describe the existing tracking approaches, including the segment-based approach. Finally, we compare the approaches and motivate the paper’s direction.

2.1 Tracking Scenario

We assume that moving objects are constrained by a road network and that they are capable of obtaining their positions from an associated GPS receiver. Moving objects, also termed clients, send their location information to a central database, also termed the server, via a wireless communication network. We assume that disconnects between client and server are dealt with by other mechanism in the network than the tracking policies we consider. When a disconnect occurs, these mechanisms notify the server, which may then take appropriate action.

After each update from a moving object, the database informs the moving object of the representation it will use for the object's position. The moving object is then always aware of where the server thinks it is located. The moving object issues an update when the predicted position deviates by some threshold from the real position obtained from the GPS receiver.

Figure 1 presents a UML activity diagram for the update scenario (activity diagrams model activities that change object states). The client initially obtains its location information from the GPS receiver. It

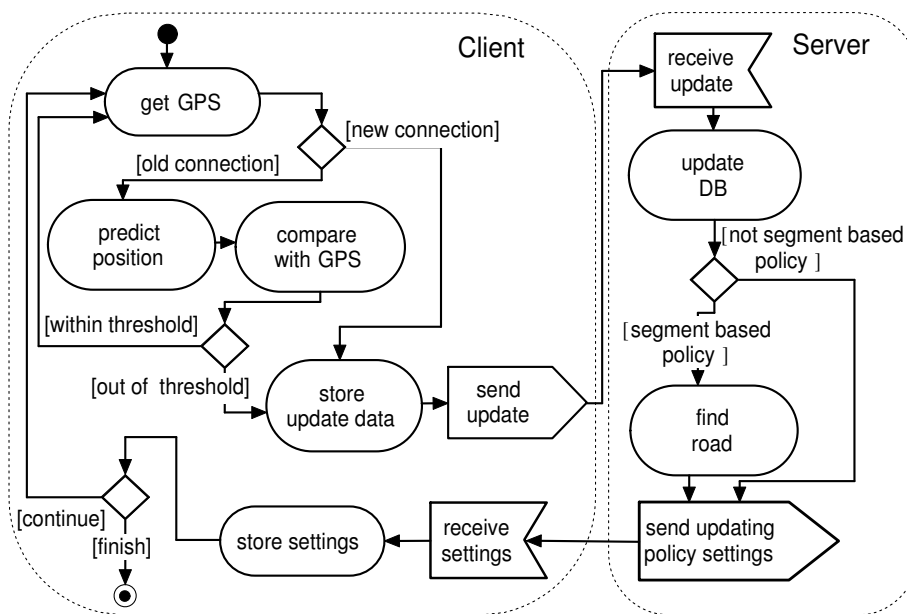


Figure 1: Tracking Scenario Diagram

then establishes a connection with the server and issues an update, sending its GPS information and unique identifier to the server.

Having received this update, the server determines which tracking approach and threshold to use for the client (these are predefined), and it stores the information received from the client in the database. If the tracking approach is the segment-based one, the server also uses map matching to determine on which road segment the client is moving. The server then sends its representation of the client's current and future position to the client.

Having received this information from the server, the client obtains its actual, current location information from the GPS receiver. The client then calculates its predicted position using the representation received from the server, and it compares this to the GPS position. If the difference between these two exceeds the given 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. Although the server may also initiate and terminate the tracking, we assume for simplicity that the client is in control. This aspect has no impact on the paper's contribution.

2.2 Data Description

As mentioned, GPS is used for positioning of the moving objects. In experiments that will be reported throughout the paper, we use GPS-log data collected during an intelligent speed adaptation project [9]. In this project, GPS receivers and small custom made computers were installed in a number of cars that were driving in the Aalborg area, Denmark. This resulted in the collection of a GPS-log for each car that contains position samples for approximately every second during the periods when the car was being operated during a period of approximately 8 weeks.

For our experiments, we also use a digital road network obtained from the same project. The road network is composed of a set of segments, each of which corresponds to some part of the road network that is in-between a pair of consecutive intersections or an intersection and a dead end. A segment is defined as a sequence of coordinates, i.e., as a polyline. Further, the road network is partitioned into streets, and each segment belongs to precisely one street. Each segment identifies its street by referring to a street code. The top part of Figure 5 offers a visual image of part of the digital road network.

2.3 Existing Tracking Approaches

We proceed to describe three existing tracking approaches [4]. With minor variations, the first and third of these were previously proposed by Wolfson and Yin [18] (see Section 6 for additional discussion). These techniques follow the scenario described in Section 2.1, but differ in how they predict the future positions of a moving object.

Point-Based Tracking. Using this approach, the server represents a moving object’s future positions as the most recently reported position. An update is issued by a moving object when its distance to the previously reported position deviates from its current GPS position by the specified threshold. An example of point tracking is presented in Figure 2(a). Here, the circles indicate the threshold, and (solid) points indicate (server-side) predicted positions that result from an update being issued by the object. The two bold lines indicate connected segments of the road network, and the thin line represents the actual object movement.




Figure 2: Tracking Policies: (a) Point-Based, (b) Vector-Based, and (c) Segment-Based

Vector-Based Tracking. In vector tracking, the future positions of a moving object are given by a linear function of time, i.e., by a start position and a velocity vector. Point tracking corresponds to the special case where the velocity vector is the zero-vector.

A GPS receiver computes both speed and heading for the object it is associated with—the velocity vector used in this representation is computed from these two. Using the same notation as the previous

figure, Figure 2(b) shows also the velocity vectors that are used for prediction. Solid points again indicate predicted positions that result from updates, while the remaining positions are simply predicted.

Segment-Based Tracking. Here, the main idea is to utilize knowledge of the road network in which the clients are moving. A digital representation of the road network is required to be available. The server uses the GPS location information it receives from a client to locate the client within the road network. This is done by means of map matching, which is a technique that positions an object on a road-network segment, at some distance from the start of that segment, based on location information from a GPS device.

In segment-based tracking, the future positions of a client are given by movement at constant speed along the identified segment, which is represented as a polyline. The speed used is the speed most recently reported by the client. When or if the predicted position reaches the end of its segment, the predicted position remains at the end from then on. In effect, the segment-based tracking switches to point-based tracking.

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 moving object 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.

Map matching may fail to identify a segment for several reasons. For example, the map available may be inaccurate, or it may not cover the area in which the client is located. The use of vector-based tracking within the segment-based tracking renders segment-based tracking robust.

An example of segment-based tracking is shown in Figure 2(c). Notice that all predicted positions are now located on the road segment, not on the trajectory obtained via the GPS receiver. The example next further explains segment-based tracking.

Example 2.1. Consider a taxi moving in a road network. The taxi starts a trip at 2 p.m. It starts at position (x_0, y_0) , and it travels at 50 km/h. The threshold is 100 m, i.e., we need to know where the taxi is within 100 m. With point-based tracking, an update is issued when the taxi gets to be more than 100 m away from the previously reported position. This situation is shown in Figure 2(a). Using the vector-based approach, the taxi's movement direction is taken into account. This yields a better approximation of the taxi's movement, thus reducing the number of updates—see Figure 2(b). If we have available a digital representation of the road network in which the taxi is moving, segment-based tracking is possible (see Figure 2(c)). Here, the updated point (a) occurs because the taxi slows down so that its predicted position moves ahead of the real position by more than 100 m. The updated point (b) occurs because the taxi reaches the end of a segment so that its predicted position stops, while the taxi keeps moving. The updated point (c) places the taxi on a new road segment.

2.4 Comparison of Update Policies

The tracking approaches described in Section 2.3 were evaluated by Čivilis et al. [4] using the INFATI data described in Section 2.2. Approximately 458,000 GPS positions collected from four cars were used, and thresholds ranging from 40 to 1,000 m were investigated.

Experimental results are presented in Figure 3. The results were obtained by simulating the scenario described in Section 2.1. Specifically, the movement of each car was simulated using the log of GPS positions for the car. So a client program and a server program interact, and a simple experiment management system is in charge of the bookkeeping needed in order to obtain the performance results. Instead of obtaining GPS positions from a GPS device in real time, the client program utilizes the GPS log, which of course makes the simulation much faster than the reality being simulated. The bookkeeping involves the counting of updates sent from the client program to the server program and keeping track of time.

All performance studies reported in this paper follow this pattern. The studies differ in the specific GPS data and road networks used, and in the tracking policies used.




Figure 3: Comparison of Update Policies

In Figure 3, accuracy threshold values in meters are on the x axis. The client obtains a GPS position from the GPS device every second and performs a comparison between the GPS position and the predicted position. The y axis then gives the average number of seconds in-between consecutive updates sent from the client to the server in order to maintain the required accuracy.

It is seen that the time in-between updates increases as the accuracy threshold increases, i.e., as the required accuracy decreases. The point policy shows the worst performance. Notice that the largest improvement of the segment-based and vector policies over the point policy is for smaller thresholds, while for larger thresholds the improvement is smaller. For thresholds below 200 m, the segment-based and vector policies are more than two times better than the point policy.

We find that segment-based tracking was outperformed because the road segment in the underlying road network were relatively short, having an average length of 174 m. It may be that a relatively straight road is represented by several segments. In this case, vector-based tracking may need less updates. So, although vector-based tracking is simpler and slightly better, we find it likely that it is possible to improve the segment-based tracking to be the best.

In addition, segment-based tracking, by relating the location of a moving object to the underlying road network, offers additional advantages:

- Buildings, parking places, traffic jams, points of interest, traffic signs and other road-related information that is mapped to the road network can be associated with the location of a moving object.
- Road network-based distances can be used in place of Euclidean distances.
- Acceleration profiles, driver behavior on crossroads, and other road-related data that increase the knowledge about the future positions of moving objects can be exploited.

Consequently, we have chosen to base our proposals for new and more efficient tracking techniques on the segment-based tracking approach.

3 Modification of the Road Network

Recall that with segment-based tracking, the predicted position of an object moves at constant speed along a segment until it reaches the end of the segment, at which time the predicted position remains at the end of the segment. The experimental study reported in the previous section indicates that the numbers of updates in segment-based tracking are closely correlated with the numbers of changes of segments. This motivates modification of the underlying road network representation that may lead to less segment changes.

We proceed to present several road network modifications. The main idea is to connect the road segments in such a way that moving objects would have to change segments as few times as possible as they travel in the road network. We first present a general segment connection algorithm and road network modification approach. Then three subsequent algorithms are presented that reuse this algorithm. At the end, the effects on tracking of the three algorithms are compared experimentally, and city and suburban driving are compared.

3.1 General Segment Connection Algorithm

The general segment connection algorithm **GSC** captures the overall approach to road network modification.

The idea is to iterate through all segments in the road network to be modified according to some specified ordering. During each iteration, the algorithm thus orders all available segments and then tries to extend the topmost, or current, segment with other segments. To do this, the algorithm identifies all existing segments that start or end at the start or end of the current segment and extends the current segment with the most attractive such segment(s) according to some other specified ordering. A current segment that has been extended is considered in the next iteration, but the segment(s) that were used for the extension are disregarded. A current segment that has not been extended becomes part of the result and is not considered any further.

The algorithm takes four parameters as input. The first is a road network, denoted by m , which is a set of segments. Each segment is a polyline that represents a small, linear part of the road network. Segments can have connections with other segments only at their start and end points. Further, each segment (initially) belongs to only one street and has one street code assigned to it. Additional detail about the concrete road network used in empirical evaluations in this paper can be found elsewhere [9]. The second parameter of **GSC** is a Boolean valued variable stc that controls the segment connection procedure by allowing or disallowing the connection of polylines with different street codes. The third and fourth parameters, $rnPrioritization$ and $candPrioritization$, are sort order specifications that specify how to sort sets of polylines. By supplying algorithm **GSC** with different parameters, different transformations of a road network result.

Algorithm **GSC** uses function $\text{first}(set_of_segments, stc, Prioritization)$. This function returns the segment in $set_of_segments$ that is first according to $Prioritization$, which is a sort order specification. It consists of a list of segment properties, e.g., length, speed limit, number of neighboring segments, along with an indication of whether sorting should be done in ascending or descending order.

A property of a segment can be calculated based on the other segments available in the argument set of segments. An example is the number of segments with which a segment can be extended. When calculating such properties, if stc is set to $true$, segments with street codes that are different from the street code of the current segment are not considered; otherwise, all segments are considered.

A property such as the angle between two spatially connected segments involves two segments. In this case, function `first(set_of_segments, stc, Prioritization, pl)` takes an additional parameter: segment `pl`. Then each segment from `set_of_segments` will have a property “angle” that is equal to the angle between the segment and `pl`.

Algorithm **GSC** is defined next and explained in the following.

Algorithm GSC($rn, stc, rnPrioritization, candPrioritization$)

1. $cn \leftarrow \emptyset$
2. **while** $rn \neq \emptyset$ **do**
3. $pl \leftarrow \mathbf{first}(rn, stc, rnPrioritization)$
4. $rn \leftarrow rn \setminus \{pl\}$
5. $epls \leftarrow \emptyset$
6. **for each** $pd \in \{\mathbf{start}(pl), \mathbf{end}(pl)\}$ **do**
7. $cand \leftarrow \{plc | plc \in rn \wedge (pd = \mathbf{start}(plc) \vee pd = \mathbf{end}(plc)) \wedge (plc.streetcode = pl.streetcode \vee \neg stc)\}$
8. **if** $cand \neq \emptyset$ **then**
9. $pl \leftarrow pl$ extended with $\mathbf{first}(cand, stc, candPrioritization, pl)$
10. $epls \leftarrow epls \cup \{\mathbf{first}(cand, stc, candPrioritization, pl)\}$
11. **if** $epls \neq \emptyset$ **then** $rn \leftarrow (rn \setminus epls) \cup \{pl\}$
12. **else** $cn \leftarrow cn \cup \{pl\}$
13. **return** cn

A variable cn that will accumulate the result of the algorithm is first initialized. The algorithm then iterates through the polylines of the argument road network in the argument road network in prioritization order. During each iteration, the algorithm will use up to two polylines of the road network for extending polyline pl . Variable $epls$ holds these polylines.

Lines 7 to 10 are iterated through for the two delimiting points of polyline pl . These points are returned by functions $\mathbf{start}(pl)$ and $\mathbf{end}(pl)$. Line 7 computes the set of candidate extension polylines, $cand$, for a delimiting point. If stc is true, extension polylines must have the same street code as the polyline being extended.

If candidate extension polylines exist, the algorithm proceeds with lines 9 and 10; otherwise, it proceeds with the next delimiting point or maintenance of rn and cn . In line 9, the first of the candidate polylines according to the argument candidate prioritization sort order is identified and used for extending polyline pl . The polyline used for extending pl is added to set $epls$ in line 10. Next, if extension was successful, the polylines used for extension are subtracted from rn , and the extended polyline is added to rn . Otherwise, polyline pl is added to the result set.

The algorithm returns the modified road network. It should be noted that the algorithm does not modify the street codes of segments. If stc is true, a segment is extended only with segments with identical street code, which implies that the street codes on the resulting road network are correct. If stc is false, the street codes on the resulting road network are not meaningful. Subsequent algorithms do not use street codes when this is the case.

The worst case running time complexity of the algorithm is $\mathcal{O}(n^3)$ where n is the number of polylines in the argument road network. The analysis is as follows. The main while loop executes at most n times. Within this loop, line 7 may involve iteration over all polylines in the road network. However, the worst case complexity is caused by the presence of the **first** function calls. In the worst case, it takes n iterations to determine the value for an attribute specified in a sort ordering, and it takes another n iterations to find the first element given the sort ordering attribute values. It should be noted that the road-network modification approaches based on this algorithms are executed only once and in an initial, off-line preprocessing step. The running times of the modifications thus do not affect the run-time performance of the tracking.

3.2 Street Code-Based Approach

The general idea is to give priority to connecting polylines with the same street code. This way, longer road segments are constructed that tend to correspond to parts of named streets. In cases where there are several candidates with the same street code, priority is given to the shortest polyline. This strategy reduces the probability that unconnected polylines will be short. The polyline connection algorithm for the street code-based approach is defined as follows:

Algorithm NSC(rn)

1. $rn \leftarrow \mathbf{GSC}(rn, true, [streetcode_{asc}, length_{asc}], [sides_{desc}, length_{desc}])$
2. $rn \leftarrow \mathbf{GSC}(rn, false, [streetcode_{asc}, length_{asc}], [sides_{desc}, length_{desc}])$
3. **return** rn

Algorithm NSC makes two calls to algorithm GSC using differing parameters. The first call uses the argument road network rn and requires that polylines being connected have identical street codes.

The sort order used for specifying the iteration over the road network sorts polylines primarily according to ascending street codes and secondarily according to their ascending lengths. The algorithm thus processes segments in street code order and gives priority to short segments. Notice that a sorting order of $streetcode_{desc}$ would also work instead of $streetcode_{asc}$. Both ensure that the polylines with the same street code are processed together, which is what we want to achieve.

The sort order used when selecting the best candidate polyline for extending a polyline first orders the candidate polylines in descending order according to $sides$, which has value 1 or 2 depending on how many sides to which the polyline can be extended (as it is a candidate, the value is at least 1). The secondary ordering is according to descending length. As a result, candidate segments are preferred that can be extended further, and among such candidates, the longest are preferred.

The second call to GSC is applied to the result of the first call. In contrast to the first call, street codes are not taken into account when connecting polylines. The sort orderings used are the same as those used in the first call to GSC.

3.3 Tail Disconnection Approach

The street code based approach does not distinguish between main roads and side streets. The underlying observation that motivates the tail disconnection approach is that moving object can be assumed to be moving on main roads most of the time. In this approach, we thus first connect polylines disregarding side streets, termed tails, and we only subsequently take the tails into account.

Definition 3.1. (Tails) Let $rn \subset \mathcal{PL}$ be a set of polylines. A polyline $pl \in rn$ is a tail if at least one delimiting point of pl is not connected to any delimiting point of any other polyline on rn . Tails are also termed first level tails. The i th level tails in rn are those polylines that are tails in the set obtained by subtracting all tails at lower levels than i from rn . We define $\text{Tails}(rn)$ of a road network rn as the set of pairs $(pl, level)$ of a polyline pl in rn and a level number $level$ in \mathbb{N} such that pl is a tail at level $level$.

A few comments are in order. If a road network has a purely hierarchical structure, each polyline may be a tail at some level. Polylines that belong to a circular structure in a road network, i.e., a structure where each constituent polyline is connected at both ends, are not tails. A highest tail level is assigned to all non tail polylines (e.g., $1 + \max(\{level \mid (pl, level) \in \text{Tails}(rn)\})$).

The polyline connection algorithm for this approach is defined as follows:

Algorithm TSC(rn)

1. $rn \leftarrow \mathbf{GSC}(rn, true, [streetcode_{asc}, taillevel_{desc}, length_{asc}], [taillevel_{desc}, sides_{desc}, length_{desc}])$
2. $rn \leftarrow \mathbf{GSC}(rn, false, [streetcode_{asc}, taillevel_{desc}, length_{asc}], [taillevel_{desc}, sides_{desc}, length_{desc}])$
3. **return** rn

Algorithm **TSC** has the same structure as **NSC**. In line 1, the first call to **GSC** requires that polylines being connected have identical street codes.

The sort order used for specifying the iteration over the road network sorts polylines according to ascending street code, then according to descending tail level, and finally according to ascending lengths. The sort order used when selecting a candidate polyline for extending a polyline first orders the candidate polylines in descending order according to tail level, then in descending order according to *sides*, and then according to descending length.

These sort orders ensure that non-tail polylines are connected first. Tails will be used only when no non tail-polylines are available. Using always tails with the highest levels first is also beneficial, as a polyline with tail level n can be extended with a polyline with tail level $n - 1$.

The second call to **GSC** is applied to the result of the first call. Here, connections between polylines with different street codes are allowed. The sort orderings used are the same as those used in the first call to **GSC**.

It should be noted that because algorithm **GSC** does not update tail levels, a segment being extended retains its tail level. This is exactly as intended.

3.4 Direction-Based Approach

The last approach takes into account the directions of the candidate polylines at the connection point. The idea is that moving objects are expected to be moving as directly as possible towards their destinations, which means that they will tend to move as straight as possible and by making as few turns as possible.

This approach thus gives preference to polylines that continue in the same direction as much as possible when extending a polyline. Put differently, preference is given to polylines with a direction at the connection point that has a small angle with respect to the direction of the polyline to be extended, again at the connection point.

The polyline connection algorithm for the direction-based approach is defined as follows:

Algorithm DSC(rn)

1. $rn \leftarrow \mathbf{GSC}(rn, true, [streetcode_{asc}, taillevel_{desc}, angleAvg_{asc}, length_{asc}], [taillevel_{desc}, sides_{desc}, angle_{asc}, length_{desc}])$
2. $rn \leftarrow \mathbf{GSC}(rn, false, [streetcode_{asc}, taillevel_{desc}, angleAvg_{asc}, length_{asc}], [taillevel_{desc}, sides_{desc}, angle_{asc}, length_{desc}])$
3. **return** rn

The algorithm extends the **TSC** algorithm by introducing the new properties *angle* and *angleAvg*. Property *angle* denotes the angle between a polyline being extended and a candidate for use in the extension. Specifically, the line segment at the connection point of the polyline to be extended is itself extended towards the candidate polyline. This extension corresponds to a straight extension of the polyline's line segment at the connection point. Property *angle* is then the angle between the extended line segment and the line segment of the candidate polyline at the connection point. See Figure 4. A small angle is thus preferable.

Next, property *angleAvg* of a polyline being extended denotes the average of the smallest *angle* values possible for both ends of the polyline. If the polyline cannot be extended to one side, an angle of 180 degrees




Figure 4: Angles Between Polylines

is used for that side. Thus, for a polyline that can be extended with three polylines to one side with angles of 34, 22, and 90 degrees, respectively, and that has no extensions on the other side, $angleAvg = (22+180)/2$. The other parameters are the same as in the **TSC** algorithm.

3.5 Comparison of Approaches

The goal of all the road modification approaches is to connect the polylines of road segments into longer polylines, so that moving objects travel on less polylines. In doing this, we assume that objects in a road network move mostly along the main roads. We proceed to evaluate the results of the road network modifications in terms of how well the constructed polylines correspond to the main roads.

All policies succeeded in connecting short polylines into longer ones. Before modification, the road network has 14,708 segments in total, and the average length of a segment is 174 m. Application of each of the three modification approaches resulted in close to 5,800 polylines and an average length of about 450 m.

Before presenting the results of general experiments with the different modified networks, we use the examples in Figure 5 to offer the reader a feel for the approaches and their differences. The figure displays part of the unmodified road network at the top, with two smaller parts, labeled *a1* and *b1*, being identified for further consideration.

The polylines created by the street-code based approach were the worst in connecting the main roads. (In residential and other areas, it is common for a main road and its side streets to have the same street code.) Parts *a2* and *b2* exemplify how the street-code based approach fails to capture the main road as a single polyline, or segment, for these two parts of the network. In Part *a2* the main road is vertical and straight; in Part *b2*, is also fairly straight, but horizontal. The number next to a short polyline is the identifier (“id” for short) of the “long” polyline to which it belongs. A bold line without an id represents an unconnected polyline.

In Part *a2*, segment 6,303 makes a loop, and segment 5,961 starts on the side road, while the main road has a single segment that was not extended by the modification procedure. In Part *b2*, segment 6,473 “ends” because of a turn to a side road.

The tail disconnection approach fixes some of the problems. In Part *b3* of the figure, a single segment now represents the main road. This is in contrast to the situation in Part *b2*, where two segments share the main road. Specifically, we see that segment 6,850 represents the entire main road on the map and that side roads are “eliminated.” However, the tail disconnection approach does not improve the situation in Part *a2*, where there are no tail polylines.

The direction-based approach is the best at assigning main roads to few polylines. This approach solves problems like those shown in Parts *a2* and *b2*. Since priority is given to straight extensions of polylines, a single polyline, with id 5,753, represents the straight part of the road—see Part *a3*. With the direction-based approach, and unlike the two other approaches, the main road is represented by one polyline, and the side roads are represented by two polylines.

In Figure 6, we present a comparison of the update performances for the segment-based policy using the non-modified road network and the road networks resulting from application of the street code-based

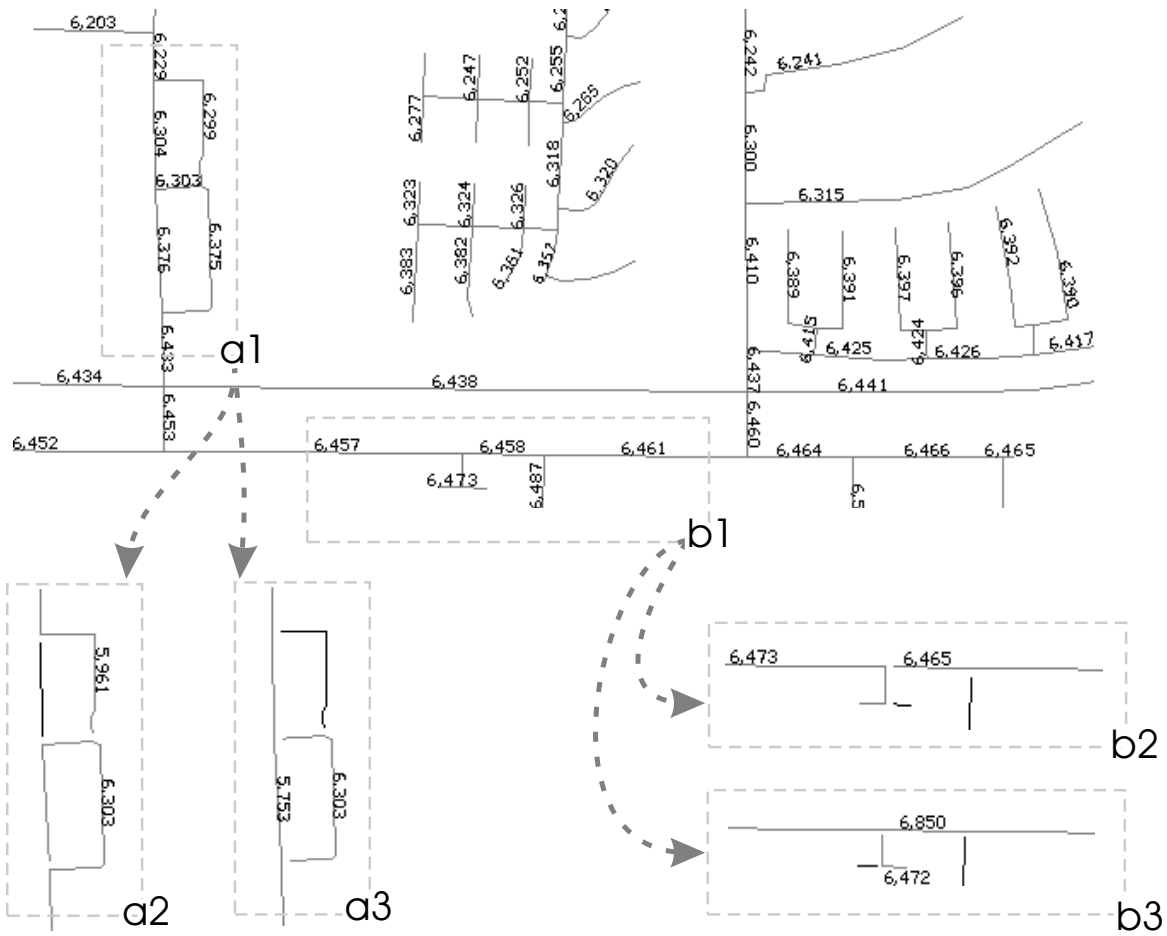


Figure 5: Road Network Modifications

approach (algorithm SSC), the tail disconnection approach (algorithm TSC), the direction-based approach (algorithm DSC). The vector-based policy is also included.

In the comparison, 568,307 GPS records were used. The curves in the upper chart show experimental results using thresholds ranging from 40 to 1,000 m, and the lower chart provide a better view of the results for thresholds in the range of 40 to 200 m.

All three road network modifications increase the performance of the segment based policy and outperform the vector-based policy. The segment-based policy has the best performance when using the road network resulting from the direction-based modification. The performance of a theoretical, constant-speed optimal policy, to be explained and discussed in the next section, is also included in Figure 6. This policy in effect assumes that a moving object always stays on the same segment and moves at constant speed, meaning that updates thus only occur due to speed variations. Based on these experiments, we select the direction-based approach as the best of the three road network modification approaches.

3.6 Comparison of Suburban Versus City Driving

Another round of experiments were conducted to see the effects of city versus suburban driving. For these experiments, we used GPS logs from 10 cars that total more than one million GPS points. We divided these points into two parts, .56 million points that are located within a rectangular region enclosing the center of Aalborg were designated as city points, and .45 million points outside this region were designated as sub-


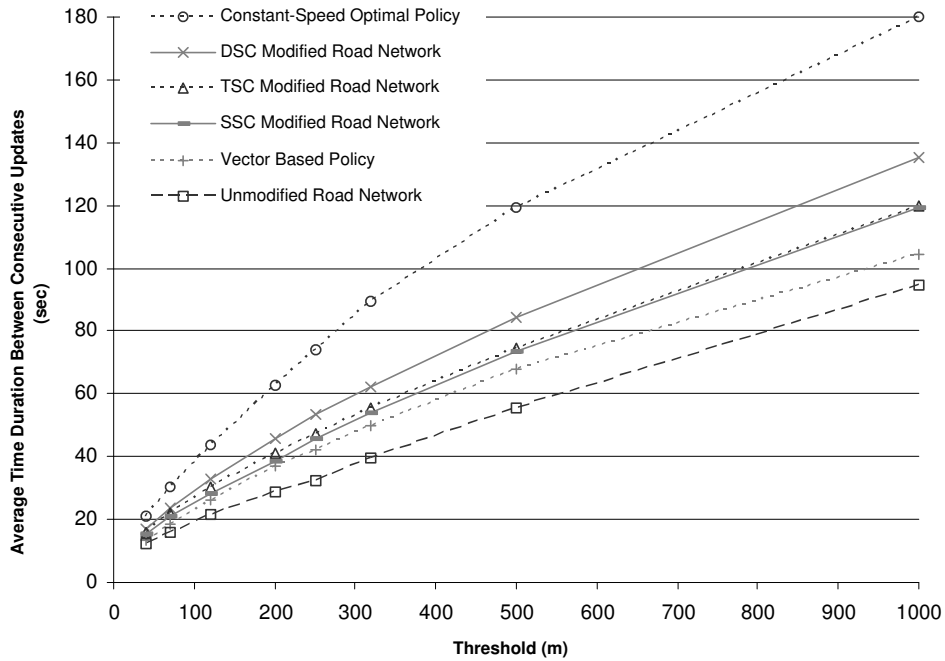


Figure 6: Comparison of Road Network Modifications

urban points. Figure 7 shows how different techniques perform for the city and suburban data. Specifically, we consider the segment-based policy with an unmodified road network and the network resulting from application of the direction-based approach transformation (algorithm DSC), and we consider the “optimal” policy. Thresholds range from 40 to 1000 m.




Figure 7: Segment-Based and “Optimal” Tracking for City and Suburban Driving

We expect smoother speed variations for suburban driving than for city driving. The better performance of the “optimal” policy (which is sensitive only to speed variations) for suburban driving than for city driving confirms this. It can also be observed that the experiments with the unmodified road network differ little for city and suburban driving. This suggests that using the unmodified road network, the majority of updates happens due to segment changes, not due to speed variation. It should also be noticed that use of the transformed road network yields better performance for city as well as suburban driving, in comparison to use of the unmodified network. This indicates that many updates caused by segment changes were avoided. Finally, it is observed that, with the modified network, the segment-based policy performs better for suburban driving than for city driving. This may be due to both the smoother suburban speed variations and longer suburban road segments.

4 Update Reduction Using Routes

The focus of this section is the use of the routes of moving objects for update reduction. At first, we introduce a theoretical, constant-speed optimal policy. Then we consider the use of a user’s routes, which are ‘long’ segments, in the segment-based policy instead of the use of road-network segments.

4.1 Theoretical, Constant-Speed Optimal Policy

One may distinguish between the updates sent from client to server based on the outcomes of the associated map matching. Recall that in segment-based tracking, when the server receives an update with a position p_i , it attempts to map match the position onto the road network, rn , to find the most probable polyline mpl and point mp on it.

With MM being the map matching function then $(mpl, mp) = MM(p_i, rn)$; and if $MM(p_i, rn) = (\text{null}, \text{null})$, the map matching is unsuccessful, and tracking is done in vector mode. Assuming that the

map matching is successful and expression $(\text{MM}(p_i, rn)).mpl$ returns polyline mpl to which a given point p_i is map matched, then if $(\text{MM}(p_{i-1}, rn)).mpl = (\text{MM}(p_i, rn)).mpl$, we say that the update is caused by *speed*, while if $(\text{MM}(p_{i-1}, rn)).mpl \neq (\text{MM}(p_i, rn)).mpl$, we say that the update is caused by a *segment change* (position p_{i-1} is that of the previous update).

The theoretical, constant-speed optimal policy introduced here indicates how few updates it is possible to achieve with the segment-based policy in the best case that occurs when a moving object travels on only one segment and no updates occur due to segment change. The policy is optimal under the assumption that the speed of a moving object is modeled as being constant in-between updates.

This policy is included here because it gives a measure of optimality under the assumption of constant-speed prediction. The policy is used for comparison purposes only and is not a practical policy. The policy is impractical because it assumes that the entire polyline along which a vehicle will ever move is known in advance. We are able to use this policy here because we have the entire GPS logs for each vehicle. Using these, we simply construct (very long) polylines that precisely track each vehicle “ahead of time.” In practice, we receive GPS positions in real time.

In Figure 6, the curve for the constant-speed optimal policy gives the lower bound for the number of updates needed by the segment-based policy. The deviation of the segment-based policy using the non-modified road network from the optimal case is substantial. Using the modified road networks, the performance is significantly closer to the optimal case. For example, for a threshold of 200 m, the use of the road network modified using the direction-based approach increases the average time duration in-between consecutive updates from 28 to 46 s in comparison to the use of the unmodified road network.

4.2 Use of Routes

It seems reasonable to assume that individuals who travel are traveling in order to reach a destination. Folklore also has it that travelers frequently use the same routes to their destinations. For example, a person going from home to work may be expected to frequently use the same route. This general type of behavior is confirmed by the GPS logs we have available [9].

Taking advantage of knowledge of the routes used by a moving object can reduce the number of updates caused by segment changes. Since a route is a sequence of partial road segments, a route is represented simply as a polyline. Therefore, the segment-based policy, which is applicable to any polylines, is also *directly* applicable to routes. All that is needed is to collect the routes of each user [2].

When using the segment-based policy with routes, we effectively assume that we know the future positions of an object. This is like in the theoretical, constant-speed optimal policy. The differences from that policy are that the polylines that represent routes are created from the road network, not from GPS logs, and that deviations from the assumed route are handled. Specifically, if an object deviates from its route, this is treated simply as a segment change. This will most likely trigger an update, but it will not lead to failure.

For experiments, we have extracted log data that represents routes from home to work of drivers represented in our GPS data. The data set contains 56,000 log entries. Using this data, Figure 8 reports the performance of the segment-based policy when routes are used, as well as of the theoretical approach based on the same data.

The policies have practically the same performance. The small deviations between the two are only visible for higher threshold values, which is due to the small numbers of updates for these. For example, at a 50 m threshold, each policy has more than 1850 updates, which renders the difference of 96 updates invisible. At a 950 m threshold, each policy has just above 228 updates, rendering the difference of 14 updates visible. The slight deviations between the policies may be explained by the differences between the routes used by the two policies. In particular, the routes constructed from GPS points and used by the theoretical policy are slightly more detailed and thus longer than the ones used by the segment-based policy.




Figure 8: Use of Routes Versus the Theoretical Approach

The conclusion is that knowing the route of an object in advance can eliminate virtually all updates caused by segment changes and thus significantly improves the performance of the segment-based policy.

5 Update Reduction Using Acceleration Profile

Even if the future trajectory of an object is known precisely and updates caused by segment changes thus are eliminated, updates still occur due to variations in speed. The reason is that the segment-based policy assumes that objects move at constant speed—it takes an update to change the speed.

In this scenario, the modeled speed of an object moving along a road is a stair function. Figure 9 presents the variation of a car’s speed along a part of its route from home to work. The stepwise constant speed is the one used by the segment-based policy with a 70 m threshold. Each new step in the stair function is marked with a dashed line and represents an update. The density of the steps depends on the threshold—smaller thresholds yield more updates.

It is reasonable to expect that more accurate modeling of the speed variation of an object along its route, e.g., using averages of the speeds during past traversals of the route, can help better predict the future position of the object as it moves along the route. Figure 10 illustrates the speed variation of one car along part of its route from home to work. Here, the thin lines represent the speeds for 12 different traversals of the route, and the solid line represents the average speed along the route.

The figure reveals a clear pattern in how fast the car drives along the route during different traversals. The geometry of the route, the driver’s habits, and the traffic situation are probably the primary causes for this correlation. Figure 11 displays the geometry of our partial route. The figure contains distance measures that allow the reader to correlate the geometry with the patterns displayed in Figure 10. The first deceleration of the car happens in preparation for negotiating a sharp 90 degree curve. Then the car accelerates, decelerates, makes a right turn, and decelerates further as it reaches a traffic light. On green, the car accelerates along a main road where it subsequently passes through two large rotaries. It can be seen that the car reaches its highest speed on the long, straight stretch of main road and that the speed as it enters a rotary is on average higher than the speed at the traffic light. It can also be seen that the car decelerates more quickly than it accelerates. We expect this type of behavior to be typical.




Figure 9: Speed Modeling Using Constant Speed Prediction




Figure 10: Speed Pattern for 12 Traversals of a Partial Route

The clear pattern in Figure 10 indicates that tracking with better performance can be achieved by more accurate modeling of the predicted, future speed of a moving object.

We consequently create an acceleration profile for capturing the average speed variation of the movement of an object along a route. It should be noted that a profile is created for each combination of a route and object using the route. Assigning profiles to the road network that are to apply to all moving objects and for all uses of the segments of the road network is expected to be less useful. We assume the presence of a




Figure 11: Geometry of Partial Route with Indications of Traffic Lights and Rotaries

separate software component that generates frequently used routes for the moving objects being tracked [2]. Having this be a separate component is reasonable, as routes are useful for other tasks than tracking.

An acceleration profile consists of acceleration values together with the distance intervals during which these values apply. A profile is created by first dividing the average speed variation along the route into intervals where the acceleration changes sign (i.e., from positive to negative or vice versa). Then the average acceleration is calculated for each interval. We define an acceleration profile apf as a sequence of $n + 1$ measures m_i and n accelerations a_i , $(m_0, a_0, \dots, m_{n-1}, a_{n-1}, m_n)$. Acceleration a_i is valid in interval $[m_i, m_{i+1})$.

To see how an acceleration profile is used, assume an object moves with speed v_0 and that its current location (“measure”) along the route is m_0 distance units after the start of the route, where m_0 belongs to the interval $[m_{begin}, m_{end})$ in which the acceleration profile has acceleration value a . Then the predicted position m_{pred} and speed v_{pred} of the object within interval $[m_{begin}, m_{end})$ at time t is given by: $m_{pred} = m_0 + v_0 t + (a/2)t^2$ ($v_{pred} = v_0 + at$).

Figure 12 exemplifies speed modeling when using an acceleration profile. The figure concerns the movement of one moving object along a route. We assume that the segment-based policy with a 70 m threshold is used. In this figure, the light vertical dotted lines mark updates. To provide better insight into the behavior of the policy used, we include the deviation between the real position of the moving object and its position as predicted by the policy.

The algorithm “Predict Positions with Segment policy and Acceleration profile,” **PPSA**, extends a previously proposed algorithm [4] with the ability to modify the speed of an object according to an acceleration profile.

The algorithm takes two parameters as input, $mopa$ and t , where the first parameter is a structure with five elements: (i) A polyline, $mopa.pl$, that specifies the geometrical representation of the moving object’s route, (ii) an acceleration profile, $mopa.apf$, for speed prediction along the route, (iii) the location of the client, $mopa.m$, given as a measure value on the route, (iv) the speed, $mopa.plspd$, of the object, and (v) the time, $mopa.t$, when the location and speed were acquired. Parameter $t > mopa.t$ is the time point for which the location of the object should be calculated. The result is the coordinates of predicted location of the object at time t .




Figure 12: Speed Modeling Using Acceleration Profile

Algorithm PPSA($mopa, t$)

1. $m_{pred} \leftarrow mopa.m$
2. $v_{pred} \leftarrow mopa.plspd$
3. $t_{pred} \leftarrow t - mopa.t$
4. **while** $t_{pred} > 0$ **do**
5. $accel \leftarrow \mathbf{getAcceleration}(m_{pred}, mopa.apf)$
6. $S \leftarrow accel.end - m_{pred}$
7. $dt \leftarrow 0$
8. **if** $v_{pred}^2 + 2 \cdot accel.a \cdot S \geq 0 \wedge accel.a \neq 0$ **then**
9. $dt_1 \leftarrow (-v_{pred} + \sqrt{v_{pred}^2 + 2 \cdot accel.a \cdot S}) / accel.a$
10. $dt_2 \leftarrow (-v_{pred} - \sqrt{v_{pred}^2 + 2 \cdot accel.a \cdot S}) / accel.a$
11. $dt \leftarrow \max(\{0, \min(\{dt \mid dt \in \{dt_1, dt_2\} \wedge dt > 0\})\})$
12. **if** $dt = 0$ **then** $dt \leftarrow S / v_{pred}$
13. $accel.a \leftarrow 0$
14. **if** $t_{pred} < dt$ **then** $dt \leftarrow t_{pred}$
15. $m_{pred} \leftarrow m_{pred} + v_{pred} \cdot dt + accel.a \cdot dt^2 / 2$
16. $v_{pred} \leftarrow v_{pred} + accel.a \cdot dt$
17. $t_{pred} \leftarrow t_{pred} - dt$
18. **if** $m_{pred} \geq \mathcal{M}(mopa.pl, mopa.pl.p_{end})$ **then return** $mopa.pl.p_{end}$
19. **return** $\mathcal{M}^{-1}(mopa.pl, m_{pred})$

The algorithm first initializes temporary variables: Variables m_{pred} and v_{pred} are set to contain starting location and speed of the moving object, and variable t_{pred} initially holds the time elapsed since the time when the moving object’s location was acquired. The object’s movement should be predicted for this duration of time. In general, several acceleration intervals are passed through during this duration of time, meaning that different acceleration values should be applied during the prediction. The algorithm iteratively calculates the time required to pass through each acceleration interval and reduces the prediction time t_{pred} with this time. When the prediction time is exhausted (line 4), the loop stops, and the algorithm calculates and returns the coordinates of the predicted location.

In line 5, acceleration value a for the predicted location of the object m_{pred} and boundary point end of the acceleration interval where acceleration value a applies are retrieved and stored in $accel$; these are returned by function **getAcceleration**. In the case where m_{pred} is equal to boundary point m_i , the boundary point m_{i+1} of the next acceleration interval is returned. If there are no more acceleration intervals, an acceleration value of 0 is returned, and the boundary point is set to ∞ . Notice that m_{pred} is initially equal to the location of the object at the time of the update (line 1).

In line 6, the distance S to the end of the acceleration interval with acceleration $accel.a$ is calculated.

The time dt required for the object to reach the end of the acceleration interval (moving with acceleration $accel.a$) is calculated in lines 9–11. This time is calculated using the quadratic equation $accel.a \cdot dt^2/2 + v_{pred} \cdot dt - S = 0$. It has solutions only if $v_{pred}^2 + 2 \cdot accel.a \cdot S \geq 0$ (line 8), and only positive solutions are valid, as the meaning of the solution is time. If there are two positive solutions, the solution with the smaller value is the valid one (line 11). If the equation has no valid solution, the result dt is equal to 0. In this case, prediction using constant speed is performed (lines 12 and 13).

After the time required to reach the end of the acceleration interval is calculated, this time is compared to remaining prediction time t_{pred} . If the time left for which prediction should be done, t_{pred} , is less than time required to go distance S , then the algorithm does prediction only for time t_{pred} (line 14). Lines 15 and 16 then calculate the predicted location m_{pred} and speed v_{pred} . The prediction time is reduced in line 17, and the loop is repeated if $t_{pred} > 0$.

Finally, the coordinates corresponding to location m_{pred} are calculated and returned. This is done in lines 18 and 19. If the predicted location m_{pred} is beyond the end of the route as described by polyline $mopa.pl$ (line 18), the end point of the polyline is returned. This is done by comparing the predicted measure on the polyline with the measure of the end point $pl.p_{end}$ of the polyline. Function \mathcal{M} calculates the measure value on a given polyline of a given coordinate point. Otherwise, the coordinate point of m_{pred} is calculated with the inverse function \mathcal{M}^{-1} , which calculates the coordinate point of a given measure value on a given polyline.

Experimental results for the segment-based policy using routes and acceleration profiles are presented in Figure 13. These experiments are based on data from the movement of five cars along different routes. The GPS data set used here consists of a total of 57,202 records. The experiments shows that the use of acceleration profiles is able to improve performance. This confirms that when knowing the past acceleration pattern of an object’s movement along a route, it is possible to more accurately predict the future positions of the object along the route. For example, using a threshold of 250 m, the average time in-between updates is increased from 72 to 98 s. We note that with acceleration profiles, we outperform the previously introduced theoretical policy that was optimal only under the assumption of constant-speed prediction.

In closing, it is also worth considering a few speed modeling alternatives and some implication of our choice. In reality, the travel speed associated with a road segment varies during the day, and different drivers may well negotiate the same segment with different speeds. By associating acceleration profiles with routes that are specific to individual drivers, we capture the variation among drivers. And because the same route (e.g., from home to work or from work to home) is typically used during the same time of the day, the variation of speeds across during the day is also taken into account fairly well. Next, if significant variations exist within the observations based on which the acceleration profile of a route is constructed, it is possible




Figure 13: Results Using Acceleration Profile

to create several speed profiles, e.g., so that rush-hour and non-rush-hour profiles are available. We did not find a need for several acceleration profiles in the GPS data we have used. Finally, we feel that the alternative of associating acceleration profiles with the road network itself leads to solutions that either will be more complex or will be less accurate.

6 Related Work

When predicting the future position of an object, the notion of *trajectory* is typically used [10, 14, 15, 19], where a trajectory is defined in 3-dimensional [14] or 4-dimensional [15] space. The dimensions are a two-dimensional “geographical” space, a time dimension, and an uncertainty thresholds dimension. A point in this space indicates when an object is in a given location and what the uncertainty of the representation of the location is. 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 our techniques, the prediction of an object’s movement is done using the speed received from the object. For more accurate prediction, we introduce acceleration profiles that allow for quite accurate modeling of the speed variation along a route. An acceleration profile is a property of the combination of a physical road network and the habits of a concrete driver.

Wolfson et al. [16] propose two location update policies, termed immediate linear and delayed linear. These do not provide accuracy guarantees, as an object does not update its location when the deviation reaches some threshold. The occurrence of an update depends on the overall behavior of the deviation, estimated using a linear function, since the last update. Experiments on simulated data show that these policies are inferior to more recent policies, also by Wolfson et al. [17]. Like ours, these offer accuracy

guarantees. Unlike ours, they 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. In contrast, we accommodate objects with memory restrictions, and we consider the case where routes are not known and where map matching may not even succeed.

Lam et al. [11] present an adaptive monitoring method (AMM) that takes into consideration not only update, deviation, and uncertainty costs, but also the cost of providing incorrect results to queries, during the process of determining when to issue updates. In AMM, the moving objects that fall into a query region need close monitoring, and a small update threshold is used for them. Objects not inside a query region may have big thresholds. Our algorithm allows different objects to have different thresholds and allows threshold to change dynamically.

Karimi and Liu [10] describe a technique for trajectory prediction. This technique 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, but are used for all objects, and they do not take into account past choices during the trip of an object. In contrast, we use complete routes. To calculate routes, not only the trajectory of a moving object, but the time of the trip and start and destination points are taken into account. Moreover we use speed 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 artificial data generated to resemble real movement data, they conclude that a version of the point-based tracking as discussed in Section 2.3 is outperformed by a tracking technique that resembles the segment-based tracking also discussed in Section 2.3. 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 metric is numbers of updates per distance unit. They consider neither road-network modification, the use of routes, nor acceleration profiles. Their versions of point and segment-based tracking assume that map-matching always works and fail if this is not the case. This was possible because the data used in experiments was generated to be perfectly map matched. We believe that the techniques presented in Section 2 are good representatives of the techniques presented by Wolfson and Yin. It should also be noted that Ding and Güting [5] 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 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, 12, 13] 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 consider scenarios where higher accuracy, 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.

Assuming a network of geo-stationary “presence” sensors, Goel and Imielinski [7] propose to use an MPEG-based prediction model in order to determine the current location of an object while using as little sensor battery power as possible.

Next, Fox et al. [6] 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 more detailed analysis of the proposed tracking techniques.

In contrast to all related work, this paper uses a substantial data set of real GPS logs for guiding the process of designing practical techniques for the tracking of moving objects.

7 Summary, Conclusions, and Future Work

The paper presents and empirically evaluates several techniques for the segment-based tracking of moving objects. These extend the basic segment-based tracking previously proposed [4]. The proposed techniques are robust generally applicable: They function even if no underlying road network is available or if map matching is not unsuccessful, and then apply to mobile objects with even stringent memory restrictions.

The performance of basic segment-based tracking is sensitive to the segmentation of the road-network representation used and to the speed variations of the moving objects. Based on these observations, the paper presents several techniques that aim to reduce the number of updates needed for segment-based tracking with accuracy guarantees:

- *Road network modification.* The segment-based representation of the underlying road network used in segment-based tracking is modified with the goal of arriving at a segmentation that enables objects to use as few segments as possible as then move in the road network. This then reduces the number of updates caused by segment changes.
- *Use of routes.* A route is a polyline, constructed from (partial) road-network segments, that captures an object's entire movement from a source to a destination. As segments are themselves polylines, segment-based tracking readily accommodates the use of routes. Routes are specific to individual moving objects, and the use of routes is expected to reduce the number of updates caused segment changes.
- *Use of acceleration profiles.* An acceleration profile divides a route into intervals with constant acceleration and thus enables quite accurate modeling of the speed of an object as it travels along a route. The idea underlying the use of acceleration profiles is to reduce the number of updates incurred by speed variations.

Experimental performance studies using real GPS logs and a corresponding real road network representation leads to the following main conclusions:

- It is possible to improve the performance of segment-based tracking by automatic re-segmentation of the underlying road-network representation. Experiments with three re-segmentation algorithms demonstrate this as well as offer insight into which types of modification are most effective in reducing the number of updates. Experiments with city and suburban driving indicate that segment-based tracking is more efficient for the latter.
- It is indeed very attractive to use pre-computed routes for the moving objects in segment-based tracking, instead of using segments from the road-network representation. The GPS logs used confirm conventional wisdom, that mobile users are creatures of habit (or efficiency) that frequently use the same routes through the road network to reach their destinations.
- The GPS data used also reveal distinctive speed patterns for the mobile users. The experimental results show that the use of acceleration profiles increases the performance of segment-based tracking.

With acceleration profiles, tracking with a 200 m accuracy guarantee can be done with an average of one update each 77 s. This is in contrast to one update every 30 s for basic segment-based tracking.

Several promising directions for future work exist. First, it would be of interest to evaluate the costs of data transmission, in terms of actual phone-bill cost for a mobile user. Such modeling should take into account the pricing policies of mobile network operators. Second, it would be interesting to study further the creation and incremental maintenance of acceleration profiles. Self-learning techniques may be applicable. Third, a road network can be modified according to the GPS data collected from all users. This way, the connection of the road segments can be based on the use of the road network by the users. This may lead to longer segments for the majority of users, thus improving the performance of the segment-based tracking.

Acknowledgments

The authors wish to thank the handling editor for careful and timely handling of the manuscript. This work was supported in part by grants from ETRI, South Korea, from the Danish National Center for IT Research, and from the Nordic Academy for Advanced Study. In addition to his primary affiliation with Aalborg University, C. S. Jensen is an adjunct professor in Department of Technology, Agder University College, Norway.

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. In *Proceedings of the 12th ACM International Workshop on Geographic Information Systems*, pp. 127–136, 2004.
- [3] J. D. Chung, O. H. Paek, J. W. Lee, and K. H. Ryu. Temporal Pattern Mining of Moving Objects for Location-Based Services. In *Proceedings of the 13th International Conference on Database and Expert Systems Applications*, pp. 331–340, 2002.
- [4] A. Čivilis, C. S. Jensen, J. Nenortaite, and S. Pakalnis. Efficient Tracking of Moving Objects with Precision Guarantees. In *Proceedings of the First International Conference on Mobile and Ubiquitous Systems: Networking and Services*, pp. 164–173, 2004. An extended version is available as DB-TR-5, Department of Computer Science, Aalborg University, Denmark, <<http://www.cs.aau.dk/DBTR/DBPublications/DBTR-5.pdf>>.
- [5] Z. Ding and R. H. Güting. Managing Moving Objects on Dynamic Transportation Networks. In *Proceedings of the 16th International Conference on Scientific and Statistical Database Management*, pp. 287–296, 2004.
- [6] D. Fox, J. Hightower, L. Liao, D. Schultz, and G. Borriello. Bayesian Filters for Location Estimation. *IEEE Pervasive Computing*, 2(3): 24–33, 2003.
- [7] S. Goel and T. Imielinski. Prediction-Based Monitoring in Sensor Networks: Taking Lessons from MPEG. *ACM Computer Communication Review*, 31(5), 17 pages, 2001.
- [8] H. Gowrisankar and S. Nittel. Reducing Uncertainty In Location Prediction Of Moving Objects In Road Networks. In *Proceedings of the 2nd International Conference on Geographic Information Science*, 2002, <http://www.spatial.maine.edu/~nittel/publications/giscience02_hari.pdf>.

- [9] C. S. Jensen, H. Lahrman, S. Pakalnis, and J. Runge. The INFATI Data. Aalborg University, Time-Center TR-79, 2004, <<http://www.cs.aau.dk/TimeCenter>>.
- [10] H. A. Karimi and X. Liu. A Predictive Location Model for Location-Based Services. In *Proceedings of the 11th ACM International Symposium on Advances in Geographic Information Systems*, pp. 126–133, 2003.
- [11] 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. In *Proceedings of the 7th International Conference on Database Systems for Advanced Applications*, pp. 218–225, 2001.
- [12] G. Li, K. Lam, and T. Kuo. Location Update Generation in Cellular Mobile Computing Systems. In *Proceedings of the Workshop on Parallel and Distributed Real-Time Systems*, p. 96, 2001.
- [13] Z. Naor and H. Levy. Minimizing the Wireless Cost of Tracking Mobile Users: An Adaptive Threshold Scheme. In *Proceedings of the IEEE INFOCOM Conference*, pp. 720–727, 1998.
- [14] G. Trajcevski, O. Wolfson, F. Zhang, and S. Chamberlain. The Geometry of Uncertainty in Moving Objects Databases. In *Proceedings of the 8th International Conference on Extending Database Technology*, pp. 233–250, 2002.
- [15] O. Wolfson. The Opportunities and Challenges of Location Information Management. In *Proceedings of the Workshop in Intersections of Geospatial Information and Information Technology*, 2001.
- [16] O. Wolfson, S. Chamberlain, S. Dao, L. Jiang, and G. Mendez. Cost and Imprecision in Modeling the Position of Moving Objects. In *Proceedings of the Fourteenth International Conference on Data Engineering*, pp. 588–596, 1998.
- [17] O. Wolfson, A. P. Sistla, S. Chamberlain, 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 Consumption in Tracking and Location Prediction. In *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. In *Proceedings of the Third ACM International Workshop on Data Engineering for Wireless and Mobile Access*, pp. 56–60, 2003.

8 Appendix

The appendix reports on additional experiments with the tracking techniques presented earlier. Various data sets from the INFATI data [9] are used in the experiments.

8.1 Update Rates for Driving on Highways Versus in Towns

We proceed to illustrate details of segment-based tracking for different accuracy thresholds and for highway and town driving. In particular, the following experiment distinguishes between updates caused by segments ending and other updates, and it illustrates the differences in update behavior between highways and other roads.

As the speed of a car is expected to be highest on highways, frequent updates could be expected there. However, because the speed variation may also be smaller on highways, where there are no sharp turns, intersections, or traffic lights, the opposite may well hold instead.

To illustrate the update behavior on a highway versus on town roads, we extract one trip that consists of about 1,900 GPS positions is about 22 kilometers long. During the first part of the trip, until the 18.2nd kilometer, the car was driving on a highway. At the 18.2nd kilometer, the car exited the highway and used town roads for the remainder of the trip. The experiment uses the segment-based policy and the DSC modified road network.

Figure 14 covers the updates that occurred during the trip, starting at the 5th kilometer. The y axis displays different accuracy thresholds, ranging from 40 meters to 1000 meters. The x axis displays the measure along the line composed of GPS positions belonging to the trip. Put differently, it shows the distance along the trip from the start of the trip.

Small circles (\odot) in black indicate the position of the predicted position of the car at the time of a “usual” update. If the predicted position of the care was at the end of a segment when an update occurred, the predicted position is indicated by a large, gray circle (\bigcirc). Next, a small plus (+) in black marks the GPS position at the time of a usual update, i.e., one where the predicted position is not at the end of a segment. If the predicted position is at the end of a segment when the update occurs, the GPS position at the time of update is shown by a large, gray plus (+).

To illustrate, consider the case in the figure where an update occurred between the 13th and 14th kilometer and where the accuracy threshold is 250 meters. Here, the update happened because the car slowed down: the predicted position (\odot) is in front of the GPS position (+). As another example, consider the update that happened between the 7th and 8th kilometer when using an accuracy threshold of 700 meters. Here, the predicted position (\bigcirc) reached the end of its current segment and stopped. Later, when the GPS position got 700 meters in front of the predicted position, the update occurred.

In addition to offering a concrete example of the workings of segment-based tracking, Figure 14 shows that for all thresholds, quite few updates occur on the highway. For thresholds that exceed 500 meters, updates only occur due to segment ends, which appear to be at 7.1, 8.2 and 14.4 kilometers along the trip. Note that for the 100 meter threshold, no update is associated with the segment end at kilometer 14.4. This is so because the predicted position was map matched to a different road segment during the update at 11.5 kilometers. In particular, the two movement directions along the highway are represented by different, but close and parallel, sequences of connected road segments of differing lengths.

On town roads, the update rate is much higher, especially for the smaller accuracy thresholds. Also, many of the updates are usual updates rather than updates that occur due to segment ends. This is the case because road segments are relatively long (especially for the main roads) and because the probability of exiting a segment at other places than the ends of the segment are significant.

Recall that the DSC modified road network was used for the experiment. For the modified road network, the predicted positions might be away from line composed from the GPS positions. This happens because




Figure 14: Highway versus Town

the predicted position continues to move on the current road segment until update is issued, even if the car has (recently) turned onto a different road segment. In Figure 14, the measure of the predicted position is the projection of the predicted position onto the line composed from the GPS positions.

8.2 Use of Routes Versus Routes With Acceleration Profiles

We proceed to look into the effects of using route with acceleration profiles versus routes without such profiles (i.e., with constant speed and zero acceleration) for tracking. We also illustrate the variation in updates during tracking for different traversals of the same route by the same driver.

Using the same type of diagram as in the previous figure, Figure 15 illustrates the locations of updates occurring while the same driver used the same route from home to work a total of 12 times. The data set for this experiment consists of about 10,000 GPS records, and the route is approximately 9 kilometers long. The set is a subset of the data set described and used in Section 5. Specifically, we use the data from one out of five cars. The data for this car is also illustrated in Figure 10. An accuracy threshold of 500 meters is used in this experiment.

In Figure 15, the y axis records the different trips, while the x axis shows records the measure along the route in meters. A circle (\circ) indicates the predicted position of the car at the time of an update, and a plus (+) indicates the GPS position at the time of an update. In Figure 15, each update is represented by a pair of a predicted position and a GPS position. The initializing update for each traversal is not displayed because it does not provide any valuable information: when this first update occurs, the predicted position is undefined, and the GPS position is projected onto the very start of the route.

Note also that the distance between the predicted position and the GPS position can exceed 500 meters—this may occur if the driver turns off the GPS receiver or if the GPS receiver does otherwise not provide any GPS positions for some time. As an example of this, consider the update in trip number 1 where the GPS position during the update was at 3600 meters, while the predicted position during the update was at 4550 meters. Here, the predicted position moved forward during a time period of 31 seconds because the client did not receive any GPS positions for this time. Such cases can be avoided by making the client send a status message to the server when the client is unable to receive GPS positions. Then the server knows that the accuracy threshold may not hold for this client for the time being.

In Figure 15, the updates that occur within the first kilometer show that the speed obtained from initial the update was consistently too low. In fact, for the updates shown first in the figure, the predicted positions are only a few hundred meters into the route, while the GPS positions are far enough ahead to cause an update. Next, it can be seen that there are very few updates from 1000 to 3500 meters. During this interval, the car was driving on a main road with no traffic lights.

Further, it can be observed that almost every trip has an update where the GPS position is at 3600 meters and the predicted position is ahead. This update thus occurred because the car slowed down. More specifically, the car reached an intersection controlled by a traffic light. The slowdown occurred in most of the trips because the driver entered the intersection from a side road and was making left turn. See Figure 11. The traffic light had different timings for the main and side roads—the light is green for 19 seconds and red for 48 seconds when approaching from the side road. Also, left turns may involve waiting for traffic moving straight through the intersection.

Figure 16 employs segment-based tracking and uses an acceleration profile. The same data and the same accuracy threshold are used for the results displayed in both Figure 15 and Figure 16.

The figures confirm that the use of an acceleration profile significantly decreases the number of updates. Using an acceleration profile, there are only few updates within the first kilometer—after the initial update, both the predicted position and the GPS position accelerate.

With the acceleration profile, there is only one update at 3600 meters, i.e., at the traffic-light controlled intersection, where the car for almost all trips had to stop and wait before turning left. Without an accelera-

tion profile, one update is consumed when the car slows down or stops, and another is consumed when the car starts accelerating upon exiting the intersection (cf. Figure 15).

The acceleration profile copes well with the two roundabouts and a traffic-light controlled intersection with a minor road in the interval from 4000 meters to 6500 meters; these cause only few updates. At 6500 meters, there is a traffic-light controlled intersection with a road of equal size. Here, here the use of the acceleration profile also leads to an improvement similar to what was observed for the intersection at 3600 meters.

In summary, for the route considered, the use of an acceleration profile increases the average time in-between two consecutive updates by almost a factor of two, from 128 seconds to 250 seconds.

Next, Figure 17 reports deviations between the speeds reported by the GPS receiver and the speeds used for prediction when updates occur. Routes of 5 cars, containing 57,000 GPS records, were used in this experiment. The x axis displays the speed deviation value obtained by subtracting the predicted speed from GPS speed at the time of update, then rounding the result off to the nearest factor of 10. The y axis displays numbers of updates. For example, Figure 17(b) shows that using the segment-based policy with routes, an acceleration profile, and an accuracy threshold of 40 meters results in more than 600 updates with a deviation between the GPS and predicted speeds between -5 and +5 kilometers per hour. In the speed deviation interval from +5 to +15 kilometers per hour, there are 260 updates.

For both the segment-based policy with routes and the segment-based policy with routes and acceleration profile, the updates become more equally distributed across different speed deviations when the accuracy threshold increases. This behavior occurs because the use of small thresholds results in updates shortly after the GPS speed and the predicted speed start to deviate.

Use of the segment based policy with routes and acceleration profiles results in smaller speed deviations than does use of the segment based policy with routes and without acceleration profiles. For example, in Figure 17(b), the biggest number of updates occurs in the speed deviation interval from -5 to +5 kilometers per hour, while in Figure 17(a) the biggest number of updates occurs in the speed deviation interval from +15 to +25 kilometers per hour. This happens because there is a smaller deviation between the predicted speed and the GPS speed when an acceleration profile is used for speed prediction (cf. Figure 12) in comparison to constant speed prediction (cf. Figure 9).

8.3 Variation in Update Rates Among Drivers

We proceed to describe how individual drivers contributed to the performance results for the tracking techniques, and how each driver used his or her car. For the experiments reported on here, we use data from 5 different cars.

Figure 18 and Figure 19 illustrate how much each car was being used during different times of the day and during the different days of the week. In Figure 18, the x axis displays the day of the week and the y axis displays the numbers of GPS records collected during a day. The figure shows that drivers were driving more during the weekdays and less during the weekends. In Figure 19, the x axis displays the hour of the day, and the y axis displays the numbers of GPS records collected during an hour. This figure shows that the majority of cars were used the most during two time intervals, with 7 to 9 hours in-between them. For example, car 4 was used most during the 5th hour and then (7 hours later) during the 12th hour; and car 1 was used most during the 6th hour and (8 hours later) during the 14th hour. This suggests that the cars collected the majority of their GPS records driving between the driver's home and work.

Figures 20 and 21 show how data from individual cars performed with different tracking policies. Data from 5 cars containing 568,000 GPS records were used in the experiments presented in Figure 20. For the results presented in Figure 21, route data for the same 5 cars, consisting of 57,000 GPS records, were used.

For each tracking policy, the different cars exhibit slightly different performance. Car 2 has the best performance for the vector-based policy. Car 4 has the best performance for the segment-based policy using

routes. Next, Car 5 does best for the point-based policy; for the segment-based policy using the unmodified road network and, to a lesser degree, the DSC modified road network; and for the segment-based policy using routes with acceleration profiles.

Although the different cars exhibit differences in tracking performance, each car has worst performance across policies for the point-based policy, and the performance increases when using the segment-based policy and unmodified road network. Use of the vector-based policy improves performance some, and use of the segment-based policy with the DSC modified road network further improves performance. Even better performance is achieved with the segment-based policy using routes and all cars have the best performance when using the segment-based policy with routes and acceleration profiles.

8.4 Geographical Distribution of Updates

We proceed to illustrate how updates from selected vehicles are distributed across the Aalborg area. Data from 5 cars, totaling some 568,000 GPS positions, are used in the experiment that also uses the segment-based policy with the DSC modified road network and a threshold of 100 meters.

Figure 22 shows the Aalborg map with an overlay of select cells from a grid with a cell size of 1 square kilometer. Specifically, all cells with at least 50 GPS are depicted. The numbers at the centers of the cells indicate the percentage of GPS positions in the cells that triggered an update. For example, 8 % of all the GPS positions in cell G11 triggered an update. The higher than average updates rate in G11 may be due to the tunnel that goes under the Lim Fjord, which crosses through the Aalborg area.

Figure 22 shows that the update rate tends to be relatively high in the center of Aalborg, which, roughly, is the square delimited by cells 13G and 10D. The update rate is also relatively high in inhabited suburbs (e.g., cells A21, I17, G3, and K5) and tends to be relatively low in uninhabited suburbs (e.g., cells I7, D18, and B21).




Figure 15: Segment-Based Tracking With an Accuracy of 500 Meters During 12 Traversals of a Route




Figure 16: Segment-Based Tracking Using Acceleration Profile and an Accuracy of 500 Meters During 12 Traversals of a Route




Figure 17: Difference Between Predicted and GPS Speed at Times of Update




Figure 18: Usage of the Cars During the Week




Figure 19: Usage of the Cars During the Day

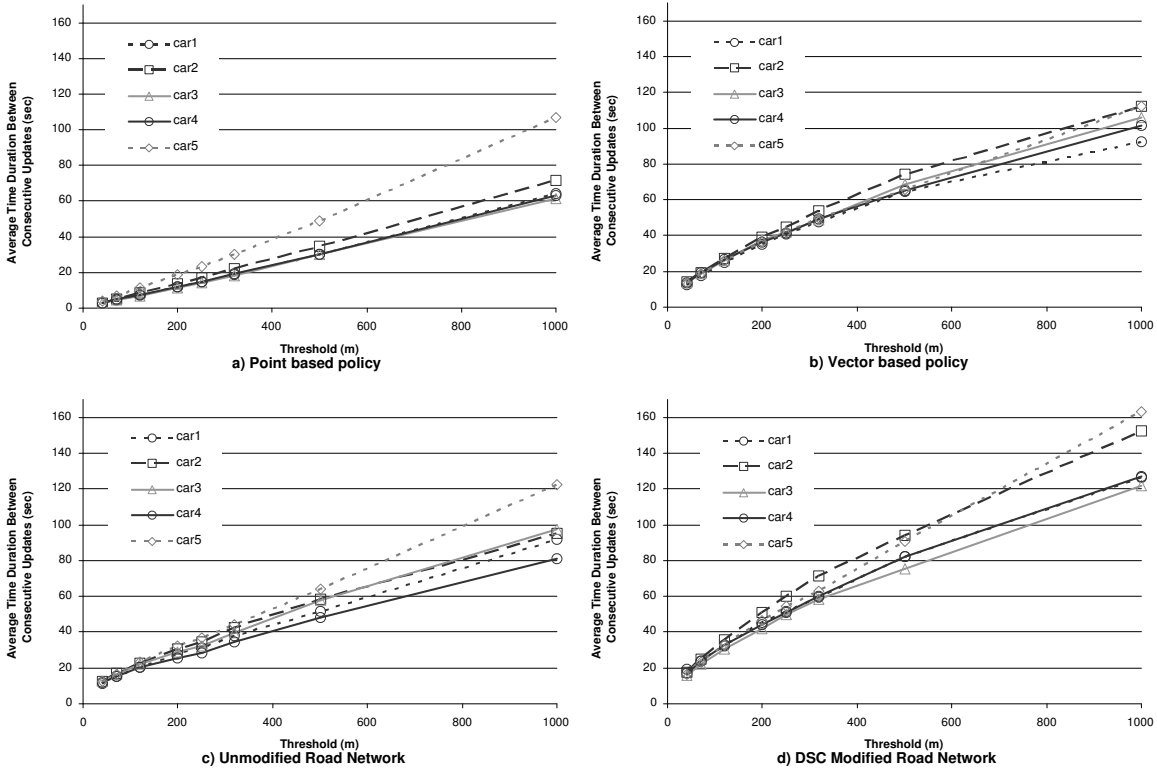


Figure 20: Performance of Tracking Policies for Different Cars

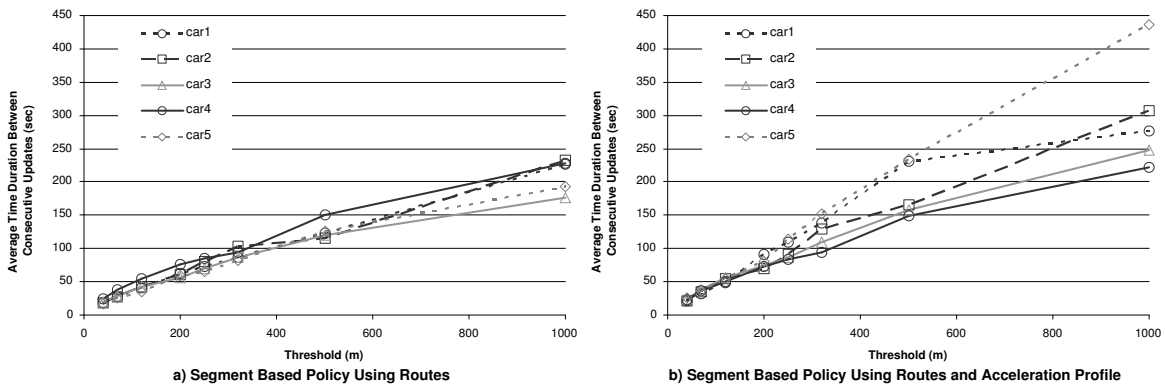


Figure 21: Performance of Tracking Policies for Different Cars Using Routes




Figure 22: Geographical Distribution of Update Frequency