

Vehicular Speed Estimation using Received Signal Strength from Mobile Phones

Gayathri Chandrasekaran*, Tam Vu* Alexander Varshavsky†, Marco Gruteser*,
Richard P. Martin*, Jie Yang‡, Yingying Chen‡

* WINLAB, Rutgers University
North Brunswick, NJ 08902
{chandra, tamvu, gruteser, rmartin}
@winlab.rutgers.edu

† AT&T Labs
Florham Park, NJ 07932
{varshavsky}
@research.att.com

‡ Stevens Institute of Technology
Hoboken, NJ 07030
{jyang, yingying.chen}
@stevens.edu

ABSTRACT

This paper introduces an algorithm that estimates the speed of a mobile phone by matching time-series signal strength data to a known signal strength trace from the same road. Knowing a mobile phone's speed is useful, for example, to estimate traffic congestion or other transportation performance metrics. The proposed algorithm can be implemented in the carrier's infrastructure with Network Measurement Reports obtained by a base station or on a mobile phone with signal strength readings obtained by the handset and depending on implementation choices, promises lower energy consumption than Global Positioning System (GPS) receivers. We evaluate the effectiveness of our algorithm on highway and arterial roads using GSM signal strength traces obtained from several phones over a one month period. The results show that the Correlation algorithm is significantly more accurate than existing techniques based on handoffs or phone localization.

Author Keywords

Received Signal Strength (RSS), Correlation

ACM Classification Keywords

I.5.4 Applications: Signal processing.

General Terms

Experimentation, Measurement, Performance

INTRODUCTION

Traffic congestion not only remains a nuisance to commuters, but also causes excess gasoline consumption, pollution, and economic inefficiency. Recent studies from the US Bureau of Transportation [2] show that nearly 4.2 billion hours were lost by drivers stuck in traffic on the nation's highways in the year of 2007. Real time traffic information can help commuters plan their commute times or help them choose alter-

nate less congested roads. Traffic information is collected from a variety of sources including inductive loop detectors embedded in the roadway [8], traffic cameras, and eyewitness reports. Recent studies [10, 14, 3] have shown, however, that it is possible to cover a larger fraction of roadways at lower cost by using mobile sensors such as GPS-equipped vehicles or smartphones. The level of traffic congestion can be estimated from the speed of these probe vehicles, as measured by the GPS receiver. A cell phone based approach is particularly attractive because it allows deploying the system more quickly to a larger number of drivers, which improves the quality of the traffic state estimation.

Current GPS-based speed measurement techniques, however, still have several drawbacks. First, frequent sampling of the GPS from the driver's smartphones leads to high battery consumption on the phones. Second, the number of programmable GPS-equipped phones is still relatively small compared to the total number of phones in use worldwide, particularly in developing regions. An alternate technique that does not face the same drawbacks, is speed estimation from cellular handoffs [9]. This technique is limited in terms of accuracy, however, since the exact handoff location can vary based on signal strength or cell load variations.

In this work, we address these challenges by estimating speed directly from signal strength profiles of mobile phones [13]. We propose a novel Correlation algorithm that is founded on the observation that the signal strength profiles along roads remain relatively stable over time; that is, multiple passes on the same road with the same speed generate similar signal strength traces. Passing over the same road at a lower or higher speed leads to either stretched or compressed version of the signal strength trace, respectively. Thus, one can determine the speed of a phone by fitting its signal strength profile onto a training signal strength profile obtained at a known speed.

The algorithm can be used with signal strength measurements obtained in the mobile handset or at the cellular base stations. In the first case, which matches our experiments setup, the mobile records signal strength measurements from the surrounding cells every second. These readings can be obtained from existing signaling messages and thus these measurements generate only negligible overhead. In the sec-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UbiComp '10, Sep 26-Sep 29, 2010, Copenhagen, Denmark.

Copyright 2010 ACM 978-1-60558-843-8/10/09...\$10.00.

ond case, base stations can collect signal measurement without requiring any modifications on the handset. This, however, is only possible when the mobile phone is active (on-going call or data connection) since the handset will then transmit Network Measurement Report messages to the associated base station. Still, this approach is promising since a large number of travelers use their phone while in a vehicle.¹ Further, the algorithm itself may run on the mobile phone or in the infrastructure, depending on system design goals.

Within this design space lie solutions that promise to substantially reduce energy consumption compared to obtaining speed readings from a handset's GPS receiver. The power draw from a GPS receiver is quite noticeable, even if the GPS is only sampled [11]. In contrast, obtaining signal strength measurements at the base station, when the phone is on call, does not create any handset energy overhead. Even when the phone is not on call, occasionally transmitting an aggregated version of the signal strength readings obtained on the handset to a server for processing (say, every 15min) would consume lower energy consumption.

Our experimental evaluation of this algorithm uses traces we collected over a period of a month on different types of roads and under different driving conditions. The main question we investigate is how the accuracy of our algorithm compares to an existing handoff-based algorithm and a localization approach that derives speed from successive estimated locations of the phone. We find that the Correlation algorithm improves accuracy by up to 42% over the handoff algorithm and by up to 71% over the localization algorithm. We also study the performance of the Correlation algorithm under varied driving conditions, on different phone models and different cellular networks (AT&T and T-Mobile) and show that it achieves consistently accurate results.

EXISTING SPEED ESTIMATION TECHNIQUES

Three existing techniques to determine a handset's speeds are estimating doppler shifts, deriving speed from location estimates, and handoff-based speed prediction. Several works [15, 16] make use of the doppler shift in frequency caused by the moving transmitter to estimate speed. [15] can only perform coarse speed classification while [16] can predict the actual speed of the mobile. These techniques require hardware support on the handset or the base station. The latter also assumes the presence of strong Line of Sight(LOS) component between the transmitter and the receiver which can make this technique impractical in a cellular environment. We therefore will concentrate in our evaluation on comparisons with the following two techniques.

Localization Algorithm

The Localization algorithm estimates the speed of a mobile phone between two points by estimating the phone's locations at the two points, calculating the distance the phone has travelled and dividing it by the time travelled [1]. In this pa-

¹Studies from Transport Research Laboratory (TRL) [5], for example, show that a quarter of British drivers use their phone while on the move.

per, we use the fingerprinting [7] algorithm for determining phone's location. The algorithm uses the RSS fingerprints obtained from 7 neighboring towers at different known locations as the training. When an RSS fingerprint is obtained from a mobile at an unknown location, the algorithm estimates the euclidean distance in signal space between this obtained fingerprint and all the training fingerprints and determines the location to be the location of the training fingerprint that yields the minimum euclidean distance.

Handoff Algorithm

The Handoff algorithm[9] involves detecting the location of the mobile based on existing knowledge of *handoff zones*. A handoff zone is the most probable location in a given road segment where the mobile switches from the current base station to a new one. Whenever a handoff occurs in the testing trace, the location of the mobile is estimated to be the location of the most probable handoff zone. A handoff typically occurs when the signal to noise ratio (SNR) drops below a certain threshold. It turns out that on any given road segment, the locations where the SNR drops below the threshold remain stable. The average speed estimate is then the distance between the previous predicted handoff location and the current predicted handoff location divided by the total time between the previous and the current handoffs.

CORRELATION ALGORITHM

The Correlation algorithm matches a measured received signal strength trace to a training trace obtained from the same road and infers the speed from the amount of stretching or compression that yields the best match. It assumes that training traces are available for the road segments subject to speed monitoring. These could be collected as part of the service provider signal measurements to determine coverage. It also assumes that the approximate starting location and the road segment the vehicle travels on is known, for example by monitoring handoff locations as shown in prior work [12].




Figure 1. Stability of RSS over time

To illustrate the insights underlying this algorithms, consider the signal traces in Figure 1. These were obtained from a mobile phone passing three times along the same road segment, twice at the same speed of 25mph and once at 50mph. Note how the two 25 mph traces are very similar. Also, note how the faster trip produced an similar RSS trace except for being compressed by a factor of two. This illustrates how despite signal variations due to fading the signal strength on average remains relatively stable at the same location, as also previously reported in [7] for stationary handsets. Thus, if we can algorithmically determine that the faster trace is a scaled version of the slower trace by a factor of 2 and we are given the speed of the slower trace, we can calculate the speed of the faster trace.




Figure 2. Generating Virtual training traces from actual training trace.

To achieve this, the Correlation algorithm creates n scaled versions of the training trace in its training phase. The number of different scaled versions n depends on the desired speed estimation accuracy. In our implementation we chose 1mph steps up to a maximum speed of 80mph, yielding $n = 80$ scaled versions of a training trace. The scaled versions are subsampled or interpolated, so that each retains a sampling rate of 1 measurement/sec. Figure 2 shows an example of creating two scaled profiles from a training profile of 25mph. Note that the total time to completion of the 25mph trace is twice as long as that for the 50 mph trace.

In the second step, its matching phase, the algorithm seeks to determine which of the n scaled versions of the training trace best fit a given measured trace, which we refer to as testing profile. To this end, it calculates the pearson's product moment correlation co-efficient [6] over the t seconds of the testing profile with t seconds of each of the scaled training profiles. The algorithm chooses the speed associated with the scaled trace that maximizes correlations.

EXPERIMENTAL EVALUATION

Our evaluation studies the accuracy of this speed estimation algorithm on a received signal strength dataset collected with severral handsets over a period of one month. We used GSM enabled HTC Typhoon phones running the Intel-POLS [4] software on the AT&T network for all our experiments. The software records the received signal strength (RSS) and cell ID from the 7 strongest cells. We used Holux GPSlim236 GPS receivers paired with the mobile phones through bluetooth for logging the ground truth location information once every second. In order to study the speed estimation accuracy of our Correlation algorithm under varied driving scenarios, we collected traces in three different environments.

1. **Highway trace:** Two experimentors collected traces while driving from home to work at varying speeds for over a month. The common route in these drives had a 14 mile highway stretch without traffic lights, which contributed a total of 36 traces.
2. **Constant Speed trace:** This experiment involved driving on a 5 mile stretch of road thrice at three different speeds: 25mph, 40mph and 55mph.
3. **Arterial roads trace:** This experiment involved 19 drives, with each drive spanning a distance of 6 miles on arterial roads with traffic lights.

Speed Estimation Accuracy

We define the speed estimation error as the difference between the estimated average speed and the one second speed calculated from the ground truth GPS locations. We study

the speed estimation accuracy of the three algorithms, Correlation, Localization, and Handoff by using the 36 highway traces first. Both Correlation and Localization algorithms use one of the traces as the training and estimate the average speed of the remaining 35 testing traces once every 100 seconds. When applying the Correlation algorithm, we introduced 100m of alignment difference between the training the testing traces to relax the assumption about having the training and the testing profiles start at the same location. The choice of 100m came from [7] where the median location estimation errors are found to be around 96m. On the other hand, the Handoff Algorithm uses 35 traces as training to estimate the handoff zones [9] and uses the median location of the estimated handoff zones to estimate the handoff position in the "testing" trace. Thus an average speed estimate is obtained everytime a handoff occurs.

Figure 3(a) presents for the three algorithms the cumulative distribution function (CDF) of the speed estimation errors calculated at one second intervals. We observed from Figure 3(a) that the Correlation algorithm achieves the best performance with a median error of 7mph, whereas Handoff and Localization have a median error of 10mph and 12mph respectively. This is very encouraging as it shows that the Correlation algorithm can significantly outperform existing algorithms.

We next study the impact of matching duration on the estimation accuracy for the Correlation algorithm. The matching duration is defined as the duration of the testing trace that is used to obtain a single average speed estimate. Choosing the right matching duration is very important because a short matching duration may not leave enough signal profile data for a good match, whereas a long matching duration will render the algorithm too inflexible to track the changing speed of a vehicle. Figure 3(b) depicts the estimation error CDF of the Correlation algorithm under different matching durations. We found that a matching duration of 200 seconds achieves the smallest error of 6.6 mph in our case.

Finally, we examine the performance of the Correlation algorithm on the dataset of the constant speed trace under different matching durations as shown in Figure 3(c). As expected, we found that the performance of the Correlation algorithm improves as the matching duration increases because more signal strength profile data is available for matching. We do not see the adverse effect of increase in the matching duration as in Figure 3(b) because in this case the vehicle speed remains constant over the entire trace.

Sensitivity to Varied Driving Conditions

We turn to examine the sensitivity of the speed estimation algorithms to different driving conditions. Figure 4 compares the performance of the Correlation algorithm with the Localization and Handoff algorithms on constant speed traces and arterial road traces.

Under constant speed traces as shown in Figure 4 (a), we observed that the Correlation algorithm outpeforms the other two algorithms. In particular, the Correlation algorithm achieves




Figure 3. Speed Estimation Accuracy comparison across three algorithms: Correlation, Localization, and Handoff

a 50% improvement on the median error over the Localization algorithm and a 67% improvement over the Handoff algorithm. This indicates that the Correlation algorithm is highly effective under both highway traces and constant speed traces.




Figure 4. Speed Estimation Accuracy of different algorithms on Constant Speed and Arterial Road traces

Figure 4 (b) presents the performance of all three algorithms for arterial road traces. The arterial road traces involved drives on roads with traffic lights, and thus the speed of the vehicles varied more over time. Because all three algorithms estimate an average speed over a matching duration, their performance degrades in the case of frequent instantaneous speed changes. However, the Correlation and Localization algorithm maintain the comparable performance with a median error of 9mph. While the Handoff algorithm underperforms the other two algorithms with a median error of 20mph.

Sensitivity to Different Networks and Phone Models.

To study the effect of a cellular network and a phone model on the performance of the Correlation algorithm, we collected three additional constant speed traces on two phones connected to the T-Mobile network and one phone connected to the AT&T network, simultaneously. The median speed estimation error was 2.3mph and 2.9mph for the phones connected to the T-Mobile network and 3.6mph for the phone connected to the AT&T network. These results show that the performance of the Correlation algorithm is consistently good across phone models and cellular networks.

CONCLUSION

We described a Correlation algorithm that exploits the stability in the signal strength profiles on any given road for vehicular speed estimation and compared its accuracy with the state of the art Localization and Handoff speed estimation algorithms. We showed that the Correlation algorithm is up

to 71% more accurate than the Localization algorithm and up to 42% more accurate than the Handoff algorithm. We also investigated the performance of the Correlation algorithm under varied driving conditions, different phone models and different cellular networks (AT&T and T-Mobile) and showed that it achieves consistently accurate results.

REFERENCES

1. AirSage Inc. <http://tinyurl.com/c782oz>.
2. Bureau of Transportation Statistics. <http://www.bts.gov>.
3. Mobile millennium. <http://traffic.berkeley.edu/>.
4. Privacy observant location system (pols). <http://pols.sourceforge.net/>.
5. Transportation Research Laboratory. <http://www.trl.co.uk/>.
6. Keith G. Calkins. E-Book, An Introduction to Statistics. <http://tinyurl.com/2jlrsv>, 2008.
7. Mike Y. Chen, Timothy Sohn, Dmitri Chmelev, Dirk Haehnel, Jeffrey Hightower, Jeff Hughes, Anthony LaMarca, Fred Potter, Ian Smith, and Alex Varshavsky. Practical metropolitan-scale positioning for gsm phones. In *UbiComp*, pages 225–242. Springer-Verlag, 2006.
8. Benjamin Coifman. Improved velocity estimation using single loop detectors. In *Transportation Research Part A*, pages 863–880, 2001.
9. D. Gundlegard and J.M. Karlsson. Handover location accuracy for travel time estimation in gsm and umts. In *IEEE ITSC*, pages 87–94, 2009.
10. Baik Hoh, Marco Gruteser, Ryan Herring, Jeff Ban, Daniel Work, Juan-Carlos Herrera, Alexandre M. Bayen, Murali Annavaram, and Quinn Jacobson. Virtual trip lines for distributed privacy-preserving traffic monitoring. In *MobiSys*, pages 15–28. ACM, 2008.
11. Kaisen Lin, Aman Kansal, Dimitrios Lymberopoulos, and Feng Zhao. Energy-accuracy trade-off for continuous mobile device location. In *MobiSys*, pages 285–298, 2010.
12. R. Sankar and L. Civil. Traffic monitoring and congestion prediction using handoffs in wireless cellular communications. In *IEEE VTC*, pages 520–524, 1997.
13. Timothy Sohn, Alex Varshavsky, Anthony LaMarca, Mike Y. Chen, Tanzeem Choudhury, Ian Smith, Sunny Consolvo, Jeffrey Hightower, William G. Griswold, and Eyal de Lara. Mobility detection using everyday gsm traces. In *UbiComp*, pages 212–224, 2006.
14. Arvind Thiagarajan, Lenin Ravindranath Sivalingam, Katrina LaCurts, Sivan Toledo, Jakob Eriksson, Samuel Madden, and Hari Balakrishnan. VTrack: Accurate, Energy-Aware Traffic Delay Estimation Using Mobile Phones. In *SenSys*, pages 85–98. ACM, 2009.
15. C. Xiao, K. D. Mann, and J. C. Olivier. Mobile speed estimation for tdma-based hierarchical cellular systems. In *IEEE VTC*, pages 2456–2460, 1999.
16. Yahong Rosa Zheng and Chengshan Xiao. Mobile speed estimation for broadband wireless communications over rician fading channels. *Trans. Wireless. Comm.*, 8(1):1–5, 2009.