



CIVIL ENGINEERING STUDIES
Illinois Center for Transportation Series No. 13-007
UILU-ENG-2013-2013
ISSN: 0197-9191

EVALUATION OF TRAFFIC-FLOW MONITORING TECHNOLOGIES: CICERO–MIDWAY SMART CORRIDOR CASE STUDY

Prepared By
Jakob Eriksson
University of Illinois at Chicago

Research Report No. FHWA-ICT-13-007

A report of the findings of
ICT-R27-91
**Evaluation of Traffic-Flow Monitoring Technologies:
Cicero–Midway Smart Corridor Case Study**

Illinois Center for Transportation

June 2013

1. Report No. FHWA-ICT-13-007		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle Evaluation of Traffic-Flow Monitoring Technologies: Cicero–Midway Smart Corridor Case Study				5. Report Date June 2013	
				6. Performing Organization Code N/A	
7. Author(s) Jakob Eriksson				8. Performing Organization Report No. ICT-13-007 UILU-ENG-2013-2013	
9. Performing Organization Name and Address Department of Computer Science University of Illinois at Chicago Room 1120 SEO (M/C 152) Chicago, IL 60607				10. Work Unit No. (TRAIS)	
				11. Contract or Grant No. R27-91	
12. Sponsoring Agency Name and Address Illinois Department of Transportation Bureau of Materials and Physical Research 126 E. Ash Street Springfield, IL 62704				13. Type of Report and Period Covered	
				14. Sponsoring Agency Code	
15. Supplementary Notes					
16. Abstract The original goal of this project was to (1) collect probe-vehicle and sensor data in the region of interest, and perform cleaning and map matching of the data; and (2) evaluate the accuracy of multiple technologies, either through direct comparison with ground truth (probe vehicles) when data are available or through simulation where direct measurement data were unavailable. Due to a substantial and unexpected lack of spatiotemporally coinciding measurement and ground-truth data, the project changed direction midway, to quickly capitalize on technology developed in the earlier stages and to attempt to remedy the underlying problem that led to our present difficulty: the unavailability of substantial and low-cost travel-time measurements. Under this new direction, the project was successful. A prototype system for low-cost traffic-flow measurement using Wi-Fi transmission monitoring has been developed and successfully demonstrated, A survey of efficient vehicle-tracking methods was performed, comparing over 100 alternative methods, leading to the development of a significantly more accurate and cost-effective method. Finally, technology developed in the early stages of the project were repurposed to facilitate the creation of the UIC Shuttle Tracker					
17. Key Words traffic, flow, monitoring, Bluetooth, Wi-Fi, GPS, probe vehicles			18. Distribution Statement No restrictions. This document is available to the public through the National Technical Information Service, Springfield, Virginia 22161		
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of Pages 10	22. Price

ACKNOWLEDGMENT, DISCLAIMER, MANUFACTURERS' NAMES

This publication is based on the results of ICT-R27-91, **Evaluation of Traffic-Flow Monitoring Technologies: Cicero–Midway Smart Corridor Case Study**. ICT-R27-91 was conducted in cooperation with the Illinois Center for Transportation; the Illinois Department of Transportation, Division of Highways; and the U.S. Department of Transportation, Federal Highway Administration.

Members of the Technical Review Panel were the following:

Jeff Galas, Chair

Abraham Emmanuel

David Zavattono

The contents of this report reflect the view of the author, who is responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the Illinois Center for Transportation, the Illinois Department of Transportation, or the Federal Highway Administration. This report does not constitute a standard, specification, or regulation.

Trademark or manufacturers' names appear in this report only because they are considered essential to the object of this document and do not constitute an endorsement of product by the Federal Highway Administration, the Illinois Department of Transportation, or the Illinois Center for Transportation.

EXECUTIVE SUMMARY

The original goal of this project was to (1) collect probe-vehicle and sensor data in the region of interest, and perform cleaning and map matching of the data; and (2) evaluate the accuracy of multiple technologies, either through direct comparison with ground truth (probe vehicles) when data are available or through simulation where direct measurement data were unavailable.

Owing to difficulties in collecting the requisite sensor and ground-truth data, the technical direction of the project was adjusted to focus on two things: the development of new technology for low-cost collection of traffic-flow measurements, and the useful re-purposing of technologies developed for the purpose of this project. The aforementioned difficulties, and new directions, are described in more detail below.

The original objective of this project was to prepare a comprehensive quantitative comparison of a number of traffic-flow monitoring technologies. An accurate comparison of this type requires the following data: (1) sensor data from several monitoring technologies, with sufficient spatiotemporal coverage along the corridor of interest (in our case, the I-55 to Midway section of Cicero Avenue), and (2) significant “ground-truth” data along the same corridor, typically in the form of high-accuracy vehicle location traces, as provided by an in-vehicle global positioning system (GPS) receiving and position-logging facility. The purpose of the “ground-truth” data is to provide an incontrovertible basis of comparison: an accurate record of individual vehicle travel times for multiple vehicles across different hours of the day, and varying travel conditions. Given overlapping data from a variety of traffic-flow monitoring technologies, and significant ground-truth data, a statistically significant, objective comparison is possible. Without such data, any comparison will by definition be subjective. Unfortunately, because of circumstances outside of the control of the principal investigator (PI) and the project’s Technical Review Panel, several alternative sources of data failed to materialize, making the desired comparison infeasible within the time frame of the project.

Once it was determined that the necessary data would not be available, the project changed direction to quickly capitalize on technology developed in the earlier stages and to attempt to remedy the underlying problem that led to our present difficulty: the unavailability of substantial and low-cost travel-time measurements.

Under this new direction, the project has been highly successful: A prototype system for low-cost traffic-flow measurement using Wi-Fi transmission monitoring has been developed and successfully demonstrated. This work was published in the proceedings of the Association for Computing Machinery (ACM) Conference on Embedded Sensor Network Systems (SenSys) in 2012.

The city of Chicago currently collects real-time vehicle GPS traces for its fleet, using a technique based on periodic sampling plus extra samples near turns. The tracking- and cost-efficiency of this method was evaluated and compared to over 100 alternative methods, leading to the development of a significantly more accurate and cost-effective method. This work is currently under submission to ACM SigMobile 11th International Conference on Mobile Systems, Applications and Services.

Moreover, map matching and other technology developed in the early stages of the project were repurposed to facilitate the creation of the UIC Shuttle Tracker, a next-generation bus tracking, route matching, and arrival time prediction system that is now deployed on the University of Illinois at Chicago (UIC) campus and publicly available online

at <http://bus.uic.edu>. UIC facilities funded the additional development required to create this service, based on work performed in the early phases of this project (ICT-R27-91). The UIC Shuttle Tracker is tremendously popular with UIC students, faculty, and staff. Our research on automatically inferring the routes, stops, and schedules of a transit system, based on data collected by the UIC Shuttle Tracker, was published in the proceedings of the ACM Conference on Embedded Sensor Network Systems (SenSys) in 2011.

CONTENTS

CHAPTER 1	TRAVEL-TIME MONITORING WITH WI-FI.....	1
CHAPTER 2	ACCURATE AND COST-EFFICIENT REAL-TIME GPS TRACKING	4
CHAPTER 3	THE UIC SHUTTLE TRACKER.....	6
CHAPTER 4	DATA COLLECTION CHALLENGES	8
CHAPTER 5	DISCUSSION AND CONCLUSIONS.....	10

CHAPTER 1 TRAVEL-TIME MONITORING WITH WI-FI

To address the cost and consequent sparsity of travel-time data, we developed and deployed a prototype solution for a novel travel-time measurement system. Our system is based on a principle of Wi-Fi-based vehicle re-identification. With today's popularity of Wi-Fi-enabled smartphones, it is reasonable to believe that a significant fraction of vehicles carry one or more such devices. When active, these phones emit Wi-Fi transmissions that contain unique identifiers, also known as MAC addresses. Given two Wi-Fi monitors actively listening for transmissions from passing vehicles, we can determine the travel time between the two monitors simply by computing the time difference between packet receptions (from the same phone) at the two monitors.

Figure 1 is a photograph of one of our Wi-Fi monitor prototypes. Here, strapped to an existing pole with a sturdy zip-tie, this low-cost outdoor device picks up transmissions from nearby vehicles and at no cost is also able to forward this information to other nearby Wi-Fi monitors through its mesh-networking capabilities. The only installation required is the zip-tie and electricity.



Figure1. Pole-mounted Wi-Fi monitor prototype. Wi-Fi monitors are based on low-cost, off-the shelf Wi-Fi base-station hardware, modified with custom traffic-monitoring firmware.

We temporarily deployed six such Wi-Fi monitors using battery packs, along the Cicero corridor as shown in Figure 2. Over a period of 90 minutes, between 11 am and 12:30 pm, each monitor collected the time stamp, MAC address, and signal strength of all Wi-Fi packets received. These data were later processed through a series of filters to remove stationary devices, etc., to produce rough speed estimates for the time period.



Figure 2. Spatial distribution of our 6-node prototype deployment along Cicero Avenue (north = left in this picture).

Figures 3 and 4 show some of the results from this prototype deployment. In Figure 3, we show travel-speed measurements for the entire stretch of the Cicero corridor. Here, a green line indicates an individual vehicle, and the blue line is the overall speed estimate in miles per hour. Reported are the mean speeds over the entire stretch, not instantaneous speeds at discrete locations. Figure 4 shows the same type of results but for a shorter stretch of road. A distinct increase in variability is discernible: This is due to the high underlying travel-time variability for a shorter road segment. Some vehicles pass the whole segment without stopping; others are less fortunate and hit one or two red lights.

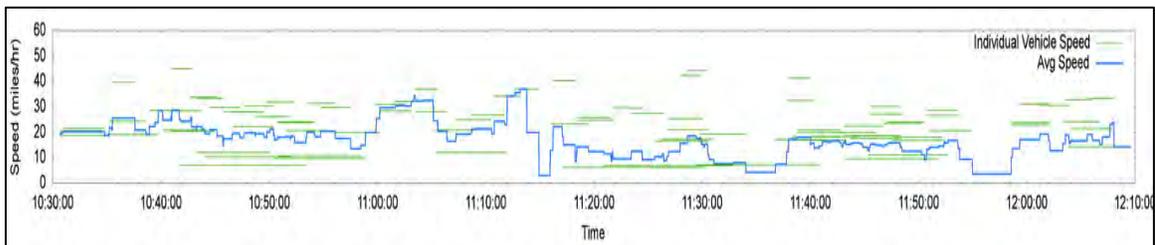


Figure 3. Travel-time measurements for 53rd to 43rd street, Cicero Avenue. Green lines indicate individual speed measurements, the blue line the average computed speed.

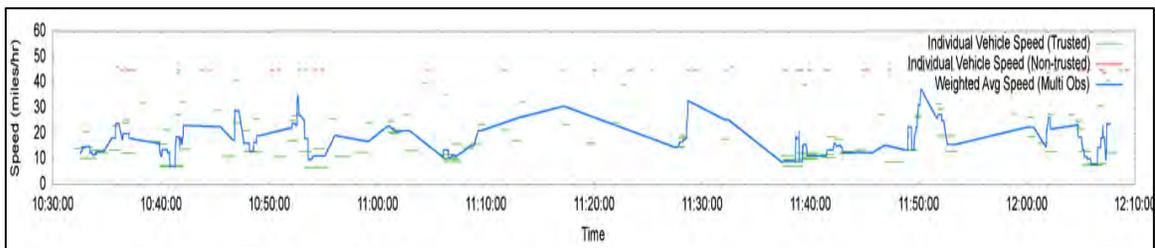


Figure 4. Single-segment travel-time measurements for 50th to 48th street, Cicero Avenue. Red lines indicate low-confidence measurements not included in mean speed.

From these results, we see that a Wi-Fi provides an effective and low-cost means of collecting travel-time data at large scale. Over 90 minutes, we collected several hundred unique travel-time measurements from passing vehicles. However, a few challenges remain: In a city, automobiles are not alone in emitting Wi-Fi transmissions—pedestrians, bicyclists, and transit riders carry smartphones as well, producing “measurement noise” that needs to be addressed. Moreover, for a larger system, questions of local speed computations, network capacity, and routing need to be addressed to ensure scalability.

Further evaluation results and in-depth technical information about our Wi-Fi vehicle-tracking system can be found in our paper “Passive Smartphone Tracking Using Wi-Fi Monitors,” published in the 9th ACM Conference on Embedded Networked Sensor Systems and available at <http://www.cs.uic.edu/Jakob/Publications>.

CHAPTER 2 ACCURATE AND COST-EFFICIENT REAL-TIME GPS TRACKING

As we describe in Chapter 4, the city of Chicago is currently using an automatic vehicle locator (AVL) system, which collects vehicle GPS traces for its fleet of more than 2,000 vehicles. Although we were unable to use this data for traffic-flow monitoring purposes because of its low granularity, it led to significant new developments in cost-efficient GPS tracking technology, as we describe below.

Current, real-time GPS tracking systems consist of a GPS receiver and a cellular network modem for transmitting the vehicle's location to a central server. Today, these systems typically report vehicle locations on a fixed time-interval basis (the Chicago Transit Authority bus tracker reportedly also sends extra samples for arrivals at bus stops). Here, a typical time interval would be between 30 seconds and 5 minutes. While more frequent reporting, such as once per second, would be desirable, it is often cost-prohibitive because of the per message charges levied by cellular providers.

The city of Chicago uses a similar system, in which the location is reported once every 5 minutes, plus an additional message for every significant turn event. Although this results in a somewhat accurate route drawing, illustrating the spatial extent of the vehicle's movements, it has extremely poor temporal accuracy (i.e., where the vehicle was along the route at any given time). This trait renders this type of data unusable for traffic-flow measurements. It also cannot be effectively used for real-time asset tracking: On average, the system knows where the vehicle was 2.5 minutes ago.

While temporal periodic sampling (such as once every 30 seconds) produces reasonably accurate tracking data, the cost of this method is unnecessarily high, especially when the vehicle is stationary. An alternative method may be to sample periodically over space, i.e., once every 300 feet. This approach may be more cost-efficient but can also result in spectacularly poor accuracy at a smaller scale: A vehicle parked 295 feet away from the most recent report would be permanently off by 295 feet. One can easily think of many different combinations of sampling policies, but it is unclear which policy is consistently better. Our most recent work to come out of this project comprehensively addresses this question and proposes a new method that significantly outperforms all known techniques, called constant-velocity extrapolation with error-based sampling. We summarize this idea below.

In contrast with current tracking systems, our proposed tracker is able to predict, or extrapolate, where the vehicle is currently located, even when no report has been received for a long time. We evaluated several extrapolation methods and found that the best accuracy was provided by a constant-velocity extrapolator. This extrapolator assumes that the vehicle continues to travel in a straight line at the velocity stated in the last report, giving the receiving server an idea of where the vehicle is at all times, not just at the instances when a location is reported. However, the extrapolated location may well be wrong, potentially exacerbating the problem.

To address this difficulty, our system includes a second extrapolator, running on the vehicle itself. This gives the vehicle a sense of where the server "thinks" the vehicle is at the moment. By sending a new location report whenever the actual location differs from the extrapolated location by a set amount, we were able to reduce dramatically the cost of real-time GPS tracking. More details on this work is available in our technical report (currently

under submission to a conference), available at <http://www.cs.uic.edu/~jakob/papers/thriftytracking.pdf>.

CHAPTER 3 THE UIC SHUTTLE TRACKER

Raw GPS coordinates are inherently noisy, as a result of physical factors such as signal reflections and atmospheric distortion. They also lack some of the information needed to produce travel-time estimates, such as what road a vehicle is currently traveling on.

The process of turning raw GPS coordinates into locations on a consistent sequence of roads traveled is called map matching. As a first-order approximation, simply “mapping” each coordinate to the nearest road segment may suffice. However, this approach is prone to significant errors, because factors such as intersections, overpasses, and adjacent parallel roads, in combination with GPS inaccuracy sometimes exceeded 50 meters. To produce accurate travel times from GPS traces, an accurate map-matched trajectory is required. As part of this project, we developed a map-matching technique based on describing the road map as a hidden Markov model (see Figure 5) and decoding the maximum probability trajectory using the Viterbi algorithm.

This technique is one of the foundations for the new UIC Shuttle Tracker developed by the PI's team from work started in this project. The UIC Shuttle Tracker consists of an AVL device in each shuttle, and a back-end server. Drivers and dispatch have no interaction with the system, which automatically determines the route currently served by each active vehicle (indicated by the color of the bus icon on the bus tracker map; see Figure 6) and continuously computes estimated arrival times for each stop.

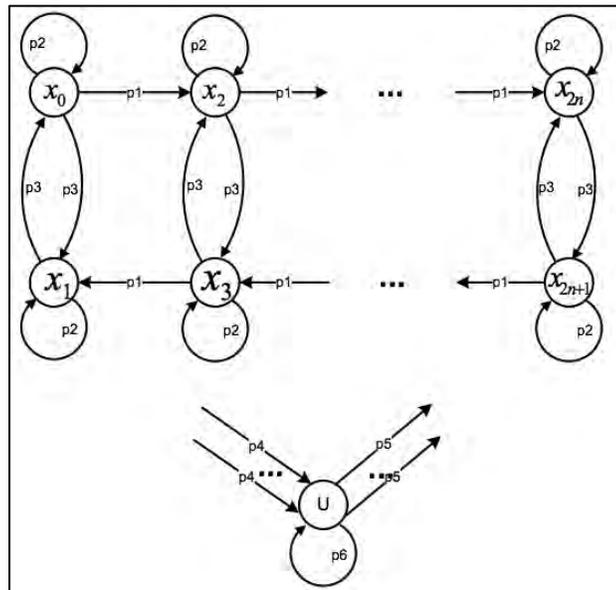


Figure 5. Hidden Markov model (HMM) describing transitions between road segments. The Viterbi algorithm is used to compute the maximum probability sequence of states for a given sequence of GPS coordinates.

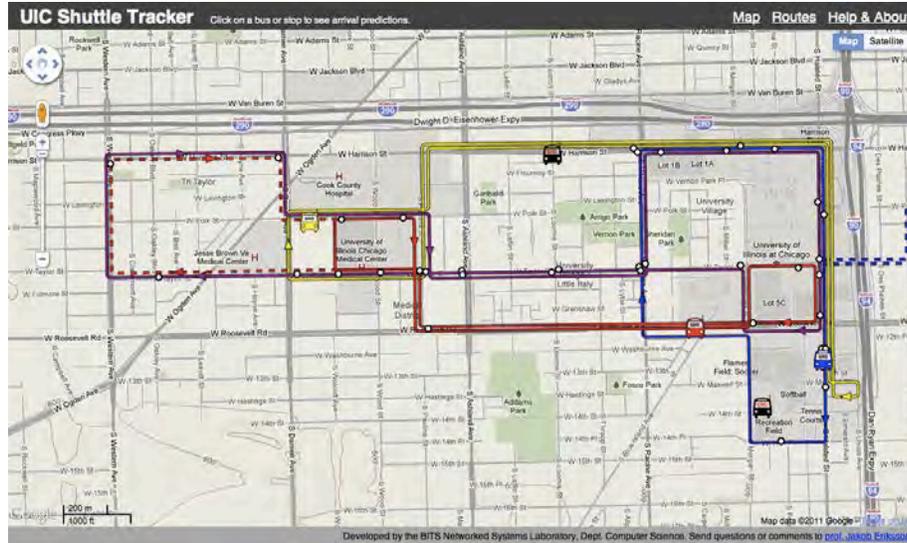


Figure 5. The UIC Shuttle Tracker is based on technology developed in the course of this project. The shuttle tracker website is used by about 1,000 unique users daily.

In addition to the interactive map, an accessible interface is offered that simply provides arrival time predictions for each route and stop, as shown in Figure 7. Automatic route classification and arrival time prediction are both derived from the map-matching techniques developed in this project.

UIC Shuttle Tracker	
Routes with real-time arrival prediction:	
UIC East Side Extended:	arrives in
Taylor & Racine	12 min
Harrison & Racine	14 min
Behavioral Sciences Building	15 min
Student Residence & Commons	15 min
Student Center East	16 min
Clinton & Cabrini	19 min
Taylor & Halsted	22 min
Roosevelt & Halsted	22 min
Beckham Hall/Stukel Towers	approaching
14th Place	8 min
note: predicted arrival times are approximate	

Figure 7. The UIC Shuttle Tracker automatically produces arrival time predictions, based on automatic service route classifications.

CHAPTER 4 DATA COLLECTION CHALLENGES

In this chapter, we briefly describe our data collection efforts as part of this project, along with the challenges encountered. This chapter should be taken as advice to future investigators that intend to pursue similar projects, rather than a report on project results.

4.1 GROUND-TRUTH DATA

The canonical ground-truth data for travel-time measurements is a periodically recorded series of time stamps and locations, for one or more vehicles. A period of 1 second is ideal; but periods up to 1 minute can be tolerated, given a sufficiently long target corridor. With data of finer granularity, small-scale behaviors, such as delays created by signalized intersections or extraneous stops along the way, can be detected and accounted for. With longer periods, only long-term average speeds are available, which can lead to lower accuracy.

For the purpose of this project, three data sources were considered for ground-truth data: Airport Express, an airport shuttle company with GPS-tracked vehicles; city taxis, which are also GPS tracked; and GPS-tracked, city-owned vehicles. Airport Express initially enthusiastically supported the project and assured us that they would make their data available to us as soon as the new technical platform was installed in their vehicles. This event was initially scheduled for June 2010, half-way into the project, which fit well with our plans. Unfortunately, owing to technical delays, these data never materialized.

In initial discussions with the panel members, taxi vehicles were thought to be a good secondary source of ground-truth data. The city works with several taxi companies that use GPS for dispatch. However, gaining access to this data turned out to be difficult, perhaps because of privacy concerns on the part of taxi companies, drivers, or passengers, though this reason was not established.

Finally, the third alternative was to make use of GPS traces from AVL-equipped city vehicles. These data were made available to the project. However, several unfortunate characteristics of these data made their use for ground-truth purposes infeasible as well. In an effort to reduce cost of service, these AVL units report vehicle locations very infrequently, as far apart as 5 minutes. To create accurate route drawings, these infrequent reports are augmented by additional location reports every time the vehicle makes a significant turn. Although this method provides good spatial accuracy, the temporal accuracy is in most cases insufficient for use as ground-truth data. Two more aspects of the data further reduced their usability as ground-truth data: (1) Being city vehicles, these vehicles often exhibit driving behaviors that are atypical for normal drivers, such as frequent stops or detours; and (2) no datum was available to identify the type of vehicle. Lacking this meta-data, a sedan used by an airport official to drive downtown may end up being treated the same as a garbage truck serving Cicero Avenue. Without significant additional processing and filtering, which in turn would have to be validated with other (currently nonexistent) ground-truth data, we were unable to make productive use of the city's AVL data for this project.

4.2 DATA SUBJECT TO COMPARISON

The original intent was to use our ground-truth data set to compare the accuracy of several technologies for travel-time measurement and estimation. These technologies

included TrafficCast BlueToad, SenSys magnetic re-identification, radio-frequency identification (RFID) toll-transponders, red-light cameras, loop detectors, and travel-time estimates derived from the Chicago Transit Authority (CTA) Shuttle Tracker. Below follows a brief description of the progress made and conclusions drawn in the attempts to acquire data from these sources.

4.2.1 TrafficCast BlueToad

Although a trial deployment of TrafficCast BlueToad sensors had been performed earlier, the collected data were neither spatially nor temporally comparable with any data collected throughout the course of this project because it came from a different road and an earlier time period. The high cost of BlueToad deployments (tens of thousands of dollars per sensor) made a new deployment infeasible.

4.2.2 SenSys Magnetic Re-Identification

Plans to perform a vendor-funded trial of this technology during the course of this project did not materialize.

4.2.3 RFID Toll-Transponders

In-vehicle toll-transponders are one commonly used technology for vehicle re-identification studies. Typically, however, these are used on highways because of the relatively high cost of installing RFID readers. No additional RFID readers were available to test this technology along Cicero Avenue.

4.2.4 Red-Light Cameras and Loop Detectors

Red-light cameras sometimes offer additional functionality supporting instantaneous speed measurements at intersections. Similarly, loop detectors, which are available in several locations along Cicero, can be used to provide volume estimates. These point-based measurements can theoretically be processed into travel-time estimates. However, at the time of our project, no end-to-end system for producing such estimates was available, meaning we would first need to develop one and train it based on ground-truth data. Owing to limited loop-sensor availability, no red-light camera availability, and essentially no ground-truth data availability, this approach was deemed infeasible.

4.2.5 CTA Bus Tracker–Derived Travel-Time Estimates

The one data source that we were able to gain full access to was an innovative bus tracker–based system developed by the Chicago Department of Transportation. Here, AVL data from city buses is processed to produce travel-time estimates for automobiles. The obvious challenge here is the difference in driving behaviors between cars and buses: Buses stop and wait for passengers and tend to be unable to take advantage of “green wave” signal timings because of frequent stops. Although this system looks promising and produces what subjectively looks to be realistic estimates, we were unable to verify its accuracy due to the lack of ground-truth data.

CHAPTER 5 DISCUSSION AND CONCLUSIONS

To summarize, lack of data availability was a substantial hindrance to performing the comparative evaluation aspects of the initially envisioned project. For a conclusive evaluation to be made, simultaneous measurements using different technologies must be conducted, on the same stretch of road over the same time interval, paired with comprehensive ground-truth measurements. We had originally thought these data were or would become available through a variety of third parties; but what data materialized was largely non-overlapping, either in space or in time, making comparison impossible.

The cost and difficulty of orchestrating this type of experiment has lead to a situation in which very little is known about the relative performance of various traffic-sensing technologies. On the basis of our experience and findings, a comprehensive comparative evaluation, without substantial (contractual) commitments from third parties, would require substantial equipment funds to be satisfactorily completed.

We would suggest that future PIs ensure not only that they have plans and backup plans for data collection but also that they ensure they have *proven access* to required data or the ability to collect said data on their own, without relying on the good will of external collaborators, before embarking on further comparison projects.

That said, the project contributed to four substantial positive outcomes: a novel Wi-Fi vehicle-tracking technology; a comprehensive evaluation of GPS tracking methods; a new error-based, real-time GPS tracking method that significantly outperforms prior work; and the UIC Shuttle Tracker.

