GPS-Based Highway Performance Monitoring Performance Monitoring Using GPS:
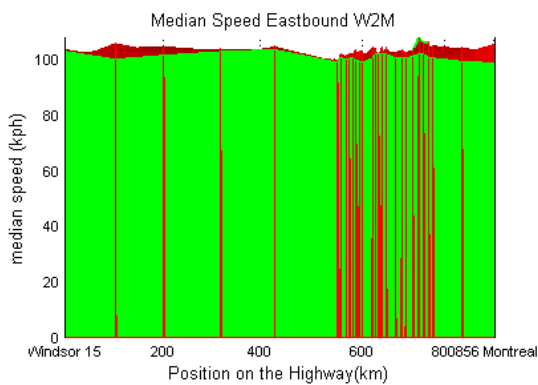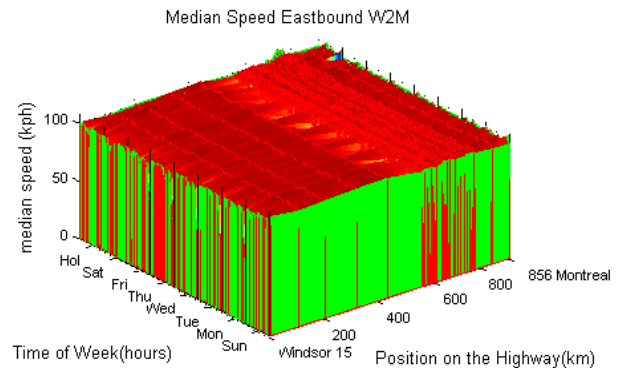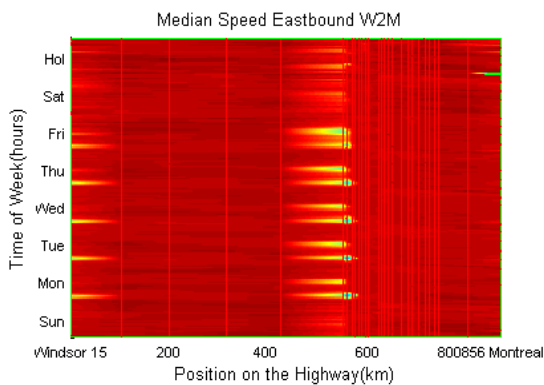
Characterization of Travel Speeds on any Roadway Segment

By

Alain L. Kornhauser, Ph. D.

Professor, Operations Research & Financial Engineering
Director, Program in Transportation, Princeton University
Also,
Founder & Board Chair, ALK Technologies, Inc.
and
Vice Chairman, New Jersey Commission on Science and Technology
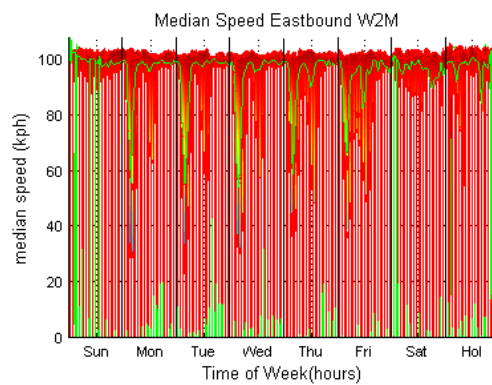
0.0 Abstract

Presented is a characterization of travel speed on any roadway segment based on probe vehicle position data. Most of the characterization is based position data obtained from GPS receivers because of their high precision and their increasing availability. Comparison is also made to Qualcomm's Automatic Satellite Position Reporting (QASPR) system because of its long history (10+ years) of extensive use by the long-haul trucking industry. Described is the use of these data in conjunction with digital map representations of roadways with particular reference to ALK's digital map database of North America. Two examples of the use of probe vehicle based GPS data to ascertain and monitor speed on roadway segments are presented. One is a demonstration of the monitoring of the speed performance of the various road segments that make up the Québec-Windsor corridor. Extensive GPS data from the first half of 2008 characterize the speed performance of the corridor by day-of-week and time-of-day. The second example also uses GPS probe vehicle data to assign a median speed , by direction, to all 31 million arcs of ALK's digital map database of North America. Examples of that assignment are displayed in geographic bandwidth charts and a generic example of a fastest route computed based on the assigned median speeds is presented.

Contents

1.0 Introduction

The quality of any roadway is characterized by the level-of-service provided by that roadway to the traveling public. While level-of-service encompasses many qualities including comfort (ride quality), beauty (scenery), grade, max load etc. arguably the most important are speed and travel time. Sensors that capture the data that characterizes the speed and travel time along a roadway segment are either embed in or near the roadway segment itself or are on-board "probe vehicles" that are traveling along that roadway segment.

Embedded sensors such as "loop detectors" and radar/laser speed detectors are placed in (loop detector) or along side of (radar/laser speed detector) the pavement to detect the presence and speed of each vehicle traveling by that specific location. The advantage of these systems is that they measure the speed of essentially every vehicle passing by that location. This produces an enormous volume of speed data that effectively characterize the temporal distribution of speed by time-of-day, day-of-week, weather and even some incidents such as police action, construction and accidents. The disadvantage is that this speed information is available only at the locations where there existed both the foresight to monitor that specific location and the substantial infrastructure funding that these systems require. Embedded systems to measure travel time require the foresight and funding to locate sensors at pair-wise locations plus the computational burden to match the data streams so as to correlate the passage of each of the paired sensors by the same vehicle. The embedded travel time systems based on matching license plate images provide travel time data on a large percentage of the vehicles that travel between the paired sensors but require a substantial computational image processing capabilities. Those based on sensing radio frequency tags (RfId), such as EZPass toll tags, enjoy a large data rate, especially in regions that have toll roads. They require substantially less data processing than the processing of license plate images. However, they also require the foresight and funding for the roadside infrastructure. Consequently, the infrastructure approaches to monitoring either speed or travel time roadway performance have been limited to the most important, but, unfortunately, infinitesimally few, locations or segments.

Vehicle-base location and speed sensors have the distinct advantage of being geographically robust in that they are not tied to a particular location and can provide speed and travel time data for essentially any roadway segment. They are limited only by physical operational constraints of the specific location and the roads traversed by these probe vehicles. The main disadvantage of these systems is that the percentage of vehicles capturing precise position, velocity, date and time in any traffic stream has been small and the percentage of those vehicles making their captured data available for roadway performance monitoring activities has been a small percentage of something that is already small. Consequently, the extent of archived position, velocity and timing data for any location is extremely small compared to the total number of vehicles that traverse any particular roadway segment.

Even though the volume of data may not be sufficient to correlate the speed and travel time performance of any segment to the level of detail that is provided by loop detector and RfId data, the geographic robustness of the data makes it imperative that these data be used to provide at least a gross level of speed and travel time monitoring over most if not all of the major highway system and at least some of the rest of the roadway system. It is especially important to develop the data architectures and computational approaches to most effectively process these data to provide the most appropriate roadway performance monitoring measures.

The price of precise vehicle-based location systems is plummeting. More individuals and businesses are installing these systems for their own benefit. Moreover, users are recognizing that these systems can be even more beneficial to them if their real-time information was shared; with managers in commercial vehicles or among a user community through integration with their cell phones' data plan in private vehicles. Thus, the ultimate beneficiaries of a geographically robust speed and travel time roadway monitoring program are the actual entities that contribute the source data that enhances the monitoring process. By contributing their location data they improve the ability to know and forecast travel times on each alternative route ahead, thus minimizing their delay and enjoying a better estimated time-of-arrival for each and every trip.

This paper focuses on the speed and travel time characterization of essentially any roadway segment using vehicle based location and velocity sensors whose outputs are tagged with a corresponding date and time. The objective of this speed and travel time monitoring system is to be the improved data source for dynamic real-time turn-by-turn route guidance systems as was demonstrated in the Albany, NY area in 2005, see Demeters, A., Kornhauser, et.al. (2006). The paper begins by reviewing the existing and needed digital map data in order to compute minimum expected time of arrival routes from anywhere at any time to anywhere. Satellite based Automatic Vehicle Location (AVL) systems are reviewed as are the characteristics of major sources of vehicle probe data. A geographically robust speed and travel time data architecture is described as are two examples of its implementation: one to demonstrate a speed monitoring program for the Québec - Montreal -Windsor Corridor and the other for maintaining median speeds on all arcs of the ALK digital map database of North America.

2.0 Digital Map Data Needs for MinETA Routing

Route guidance systems compute routes using some variant of the "shortest path" algorithm through a graph whose node and arc attributes are derived from a digital map database. The digital map database provides the source data for translating a sensor-defined stating location, typically a latitude and longitude from the US's Global Positioning System (GPS). It is currently the only operational Global Navigation Satellite Systems (GNSS). The sensor provided GPS latitude and longitude is converted to a specific location along a specific arc of the digital map database through a map-matching procedure. In the parlance of a digital map database, this point is represented by the 3-element vector: {L_o (unique arc number spanning between the location of its A node, to the location of its B node), t_o (% of distance from A node end of L) and d_o (a Boolean indicating a travel direction; A to B or B to A). The user provides the destination, often in the form of a common postal address, place name or location on a map. This user input is converted to a 3-element vector { L_d , t_d , d_d }. The shortest path algorithm then computes the route from { L_o , t_o , d_o } to { L_d , t_d , d_d } using the arc and node data contained in the digital map database. While at times a "minimum distance" route is computed where the cost on each arc is taken simply to be the distance from A to B for each arc, what is usually done is the computation of a surrogate fastest route using a modeled rough approximation of speed on each arc. Unfortunately there does not exist a digital map database that contains anything but the crudest surrogate information for travel times. Many routing algorithms, including those used by CoPilot and its ALK digital map database use a convoluted combination of distance, class of road, intersection geometry and nearby population to compute a proxy for travel times. Some use speed limit as a proxy, although there is little proof that this is a better measure of speed primarily because much of the speed limit data are inferred from road type rather than being "as posted". This is especially true for local and secondary roads. Needed are representative measures of actual speed that can be realized on each segment of the digital roadway network.

2.1 Speed and Travel Time Characteristics

Speed is a point wise performance measure at a specific location and time. Unless one makes some assumption about how speed varies between point wise values (constant or linearly varying would be the simplest two assumptions) one cannot use speed itself as a measure over anything but an infinitesimally small segment. On the other hand, travel time, **TT**, is a metric that spans a physical space at a specific time. As such it encapsulates any variation in speed that may occur over any segment and provides a summary measure of the performance of that roadway segment or arc of the digital map database.

One can convert travel time into a measured "average speed", $\underline{S}$, by dividing travel time by the distance between locations at which the travel time is measured. This normalization of travel time into a measured average speed provides a more convenient way to display, compare and understand the travel time data associated with any arc or groups of arcs in a roadway network. Since arcs and groups of arcs can have widely varying lengths, variations in travel times can be more reflective of variations in arc lengths than anything else. However, when normalized to a measured average speed one can more readily find and investigate potential data glitches and more readily monitor, compare and understand variations in road segment performances along alternative routes. For this reason, measured average speed is computed for each segment traversed by each probe vehicle.

As was shown by Kornhauser and Schrader (2004) and many others, travel times across at least some segments of the roadway network vary by time-of-day and day-of-week, as shown in Figure 1, as well as weather and nearby incidents such as construction, police activity and accidents. A proper mathematical representation of these variations may require many dimensions. Kornhauser and Schrader used a 10 parameter function to represent the variation of a measure of travel time as a function of time throughout a typical day. Speed measured at a single loop detector was assumed to remain constant throughout a road segment of known length bridging two interchanges. Thus the observed speeds were converted into a presumed travel time. (A similar function could just as easily have been defined for Speed as a function of time). This function effectively describes the temporal dynamics of up to three daily recurring congestion periods, the morning, early afternoon and evening peak congestion times, as can be seen in Figure 2.

For road segments that exhibit a strong performance variation by day-of-week as is depicted in Figure 1, one needs to estimate 10 parameters for each day for each road segment or arc. This can be for as many as 8 days, the 7 days of the week plus one for holidays, which are known to exhibit different performance characteristics. On may also wish to generate scaling factors for police activity, construction and accidents.

A great deal of data are required to characterize all of the nuances of the travel time performance of even one segment. As mentioned previously such volumes of data only exist for the precious few point locations where infrastructure speed measurement devices have been installed. Even for those, some of the exogenous data such as weather condition, police activity and construction are not necessarily well correlated or readily available with the speed detector data streams.

What is available for broad reaches of the roadway system is probe based location data. While the volume of these data is limited, it provides an observed measurement of travel time across each road segment traversed by each probe vehicle without assuming that speed is constant, linearly varying or otherwise changing across a segment. While for many segments simple assumptions are acceptable, it is known that for some segments this is a gross simplification. It can be expected that in the two examples from the interstate highway network around Milwaukee displayed in Figures 1 and 2, speeds during periods of congestion are slower near the interchange ends of segments bridging those interchanges. An even more drastic variation in speed occurs on arterials and local streets that have traffic control devices at intersections. Near the midpoints of those segments vehicles may be traveling at speeds near the speed limit, while at the ends of the segment they are accelerating from or decelerating to a complete stop. Thus any single instantaneous speed measurement will generally not properly represent the level-of-service performance of any but an infinitesimally small segment. However, travel time measures that performance directly.

Weekday Travel Time

$$TT = f(t) = K + C_1 \cdot \eta(\mu_1, \sigma_1) + C_2 \cdot \eta(\mu_2, \sigma_2) + C_3 \cdot \eta(\mu_3, \sigma_3)$$

Where :

$$\eta(\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-(t-\mu)^2/2\sigma^2}$$

and K, C_i, μ_i, σ_i *are parameters to be estimated*

(1)

Zoo Interchange – Hale Interchange (All Days)
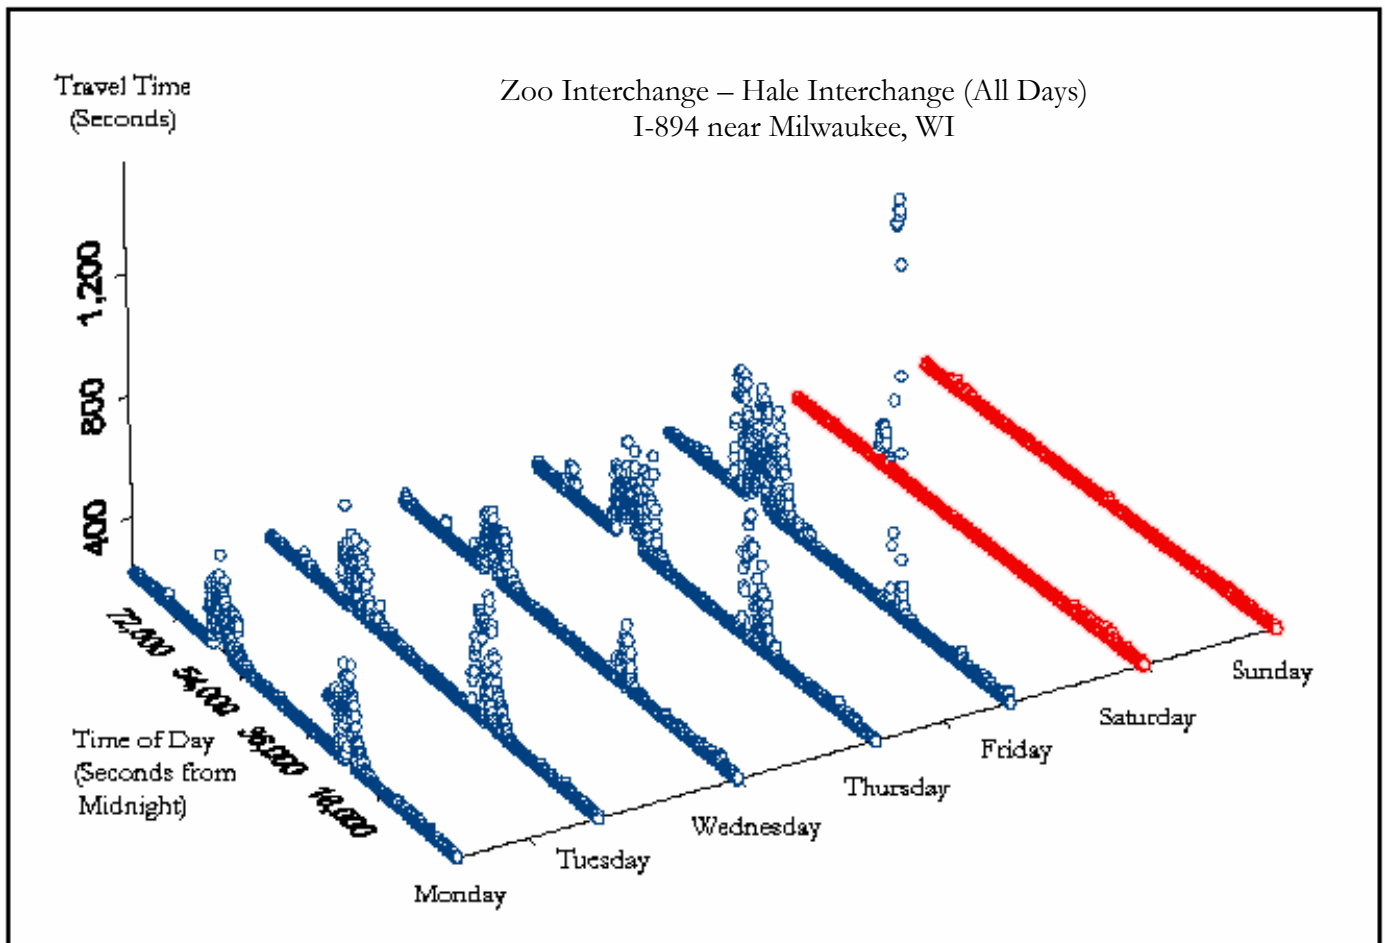I-894 near Milwaukee, WI

5

Figure 1, Example of Travel Time variations by time-of-day and day-of-week for a typical highway segment near Milwaukee, WI
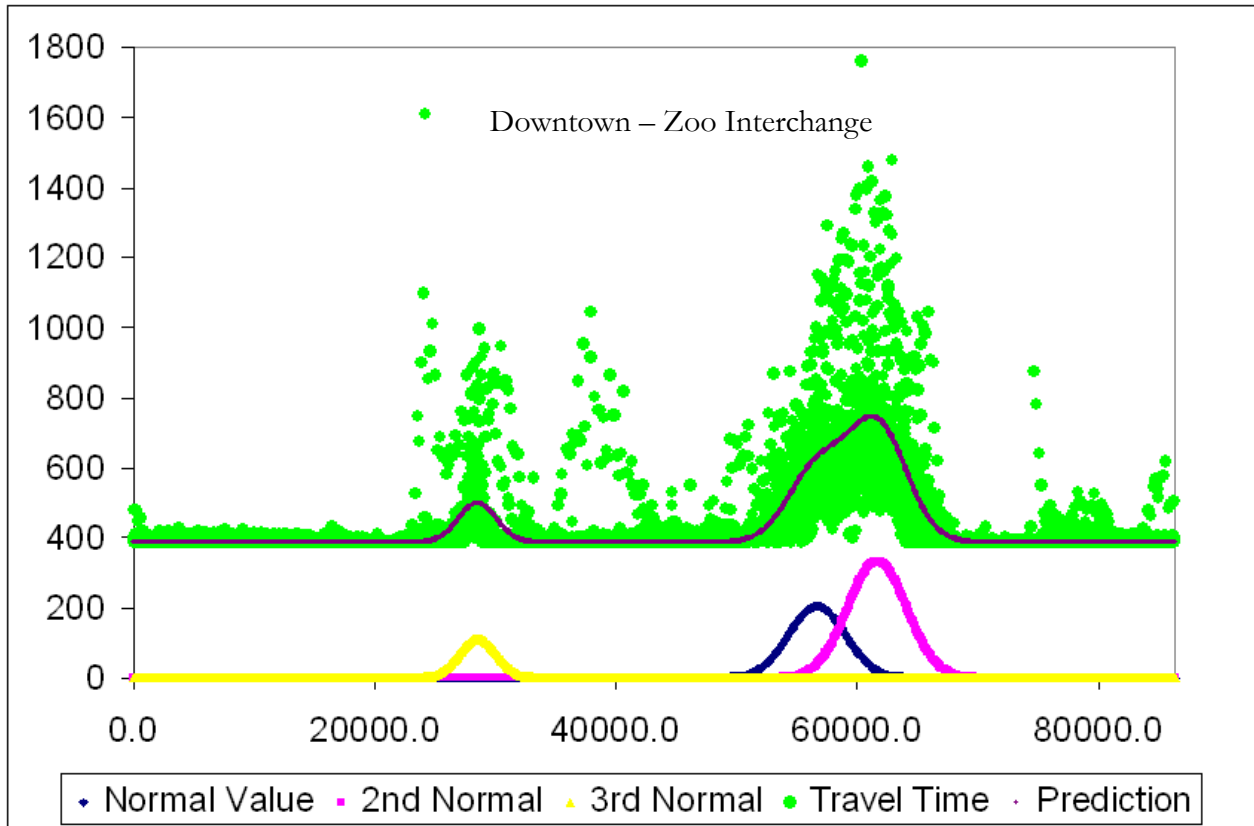

Figure 2, Elements that assemble to form the 10 parameter Travel Time function for the Downtown to Zoo roadway segment of I-94 near Milwaukee

2.2 Digital map database.

NavTeQ, TeleAtlas and ALK Technologies maintain precise geographically aligned digital map databases of the street and highway systems of North America that can be used to populate the node and arc attributes of a mathematical graph representation of the roadway system. In ALK's digital map database essentially each branch point in the roadway network is represented as a node. Each node is precisely aligned with its actual geographic coordinates. Some 31 million arcs are used to represent the centerline of the physical roadways that interconnect these branch points. Intermediate shape points are used to align the curvature of each arc with that of the actual roadway surface. The degree of precision is somewhat nebulous because the physical roadway segments are of varying width and intersections aren't actually points but extend over some area. Moreover, since the major purpose of the geographic precision is to be able to locate the position of a vehicle on the mathematical graph, there is little need to for the precision to be better than the precision of the vehicle location technology. In general, the precision of widely used GPS has a mean-squared error on the order of 10 meters which is comparable to the alignment tolerance ALK's intersections and roadway centerlines.

2.2.1 Uncertainty in the alignment of the digital map data and probe vehicle location data.

While there is substantial precision in both the vehicle location data and the digital map data, each contain error of similar distribution. Both contain high frequency positional error that is on the order of 10 meters for GPS location data and ¼ mile (400 meters) for QASPR location data. The high frequency positional error of the digital map database is of the same order, but slightly higher, as the GPS data, much of it being aligned using GPS data. The error distributions also contain a "fat tail' of low frequency error that can be substantially larger. These larger errors are primarily due to the loss of line-of-site to the GPS constellation or simply coding errors in the large digital map database. Consequently, GPS locations reported for probes traveling along any roadway will not, in general, be exactly superimposed on the digital map representation of that roadway. While usually they will be "close by" the correct roadway, the density of roads in many important areas is such that they may also be "close by" other roadways. It is the purpose of the map-matching procedure to use a "close by" metric that is able to discern among neighboring road segments, the most likely location on the digital map of each reported physical position. The map-matching process sequentially map-matches the series of location reports as they evolve to describe any trip. It uses each datum point's latitude, longitude, GPSSpeed, GPSHeading and map-match statistics from the previous datum point as input values to ALK's proprietary map-matching algorithm. This algorithm computes a goodness-of-fit index for each road segment near the GPS point. The goodness-of-fit, called "snapValue" combines, perpendicular distance, tangency, and

reachability from the previous map-matched location. The roadway location, as specified by a unique digital map segment (termed: "grid, link") and precise location along the segment (percent of segment length from the A-node end of the segment termed: "t"), having the largest snapValue is assigned as the roadway location of that GPS point. This map-matching process is applied to every GPS point. Assigned to each GPS point is its most likely roadway location on the digital map and a measure of the goodness of fit to that point.

3.0 Automatic Vehicle Location (AVL) Systems; Vehicle probe data

At present two technologies are in widespread use in to automatically locate vehicles:

- Qualcomm's Automatic Satellite Positioning System (QASPR)
  http://adsabs.harvard.edu/abs/1990imsc.conf..285A   and
- the Global Positioning System (GPS), the United States' Global Navigation Satellite System (GNSS)
  http://en.wikipedia.org/wiki/Global_Positioning_System

GPS is the more broadly used vehicle location system. In most applications it is used for personal location referencing rather than remote vehicle location monitoring. As such, until recently, it has not been normally integrated with a communications system. Possibly the most successful GPS vehicle location monitoring systems has been done by @Road's integration of a GPS receiver with cellular wireless communications for applications serving the needs of Mobile Resource Management, www.@road.com, and Qualcomm's OmniTRACS system for ubiquitous location monitoring of trucks.

3.1 Qualcomm Automatic Satellite Position Reporting (QASPR) system.

Starting in 1998, Qualcomm has utilized two geo-stationary satellites to provide both positioning and communications through its OmniTRACS mobile satellite communication system. The positioning technique uses the original OmniTRACS TDMA timing signal formats in the forward and return link directions plus an auxiliary, low power forward link signal through a second satellite to derive distance values. The distances are then converted into the mobile terminal's latitude and longitude in real time. A minor augmentation of the spread spectrum profile of the return link allows the resolution of periodic ambiguities. The system requires an unobstructed line-of-sight view of both geo-synchronous satellites. In situations where the line-of-sight is obstructed by an overpass, buildings, trees or mountains no position measurement is obtained. Otherwise, one-standard deviation positional errors of approximately 1/4 mile are consistently realized. Figure 3 is a close-up display of many of QASPR's positioning reports of trucks traveling along an unobstructed section of I-80 in Pennsylvania. All positional reportings can be assumed to have been associated with probe vehicles traveling on this section of I-81. The whisker attached to each reported vehicle location is meant to graphically represent an estimate of the instantaneous velocity of each vehicle. Speed and heading are not reported in QASPR. Only the position, date and time of anonymously identified vehicles (specifically vehicleID, latitude, longitude, date and UTC (Coordinated Universal Time)) are ever released. The data, when sorted by anonymous vehicleID, date and UTC provides a time sequenced series of points describing the vehicle's travels in space and time. This enables the tracing of individual probe vehicles and provides a convenient way to compute an approximation of the speed and heading corresponding to each reported vehicle location. The length of the whisker for the i^{th} point is proportional to the measured average speed between the i^{th} point and the $i+1^{st}$ point. This is computed by dividing the travel distance between the i^{th} and the $i+1^{st}$ point by the difference in UTC for the two point. Each whisker is drawn towards the next point. The map scale indicates that the positions are mostly contained within a ¼ mile ban on either side of the centerline of the highway with a few points having substantially larger error. Those points heading west tend to be on the north side of the centerline and those heading east tend to be on the south side of the centerline. This is a very well behaved location because this section of highway is oriented in the direction of travel for most of the probes. Since there is no other nearby road, it can be safely assumed that all of these positional reporting came from probes traveling on I-80. In more built-up areas, it is not as obvious which road was being traversed.

What QASPAR lacks in positional accuracy, reporting frequency and the lack of a corresponding velocity report is made up for by the popularity of the system by the long-haul trucking industry and its geographic ubiquity. All of the data are archived by Qualcomm and the individual trucking companies that use the system to manage their fleet. This vehicle tracking system has achieved wide acceptance in North America especially by the long-haul motor carrier industry. Given the operational nature of the long-haul trucking business, ¼ mile accuracy is sufficient for tracking purposes and there is little need for more frequent updates. Position reporting once every 45 minutes and the report of arrivals and departures tends to be sufficient to support the needs of most fleet management systems. At present there are over 300,000 trucks, mostly class 8 vehicles in long haul truck-load service, equipped with the OmniTRACS system using QASPAR. Figure 3 is a display of the instantaneous location of a one week animation of the travels of 27,417 trucks, almost a 10% sample of the QASPAR equipped fleet. Only the whiskers are being drawn. Colors are randomly assigned to vehicleIDs so that the movement of individual trucks can be identified. At this scale the animation highlights the nation's major highway corridors. Figure 4

depicts the individual movement of vehicleID #751 for the one week period. Straight lines are drawn between each position report, which are roughly 45 minutes apart. This vehicle traveled almost 5,000 miles during the 7 day period. The red vertical line in Missouri shows the location of this vehicle at 11:46 PM of the second day.
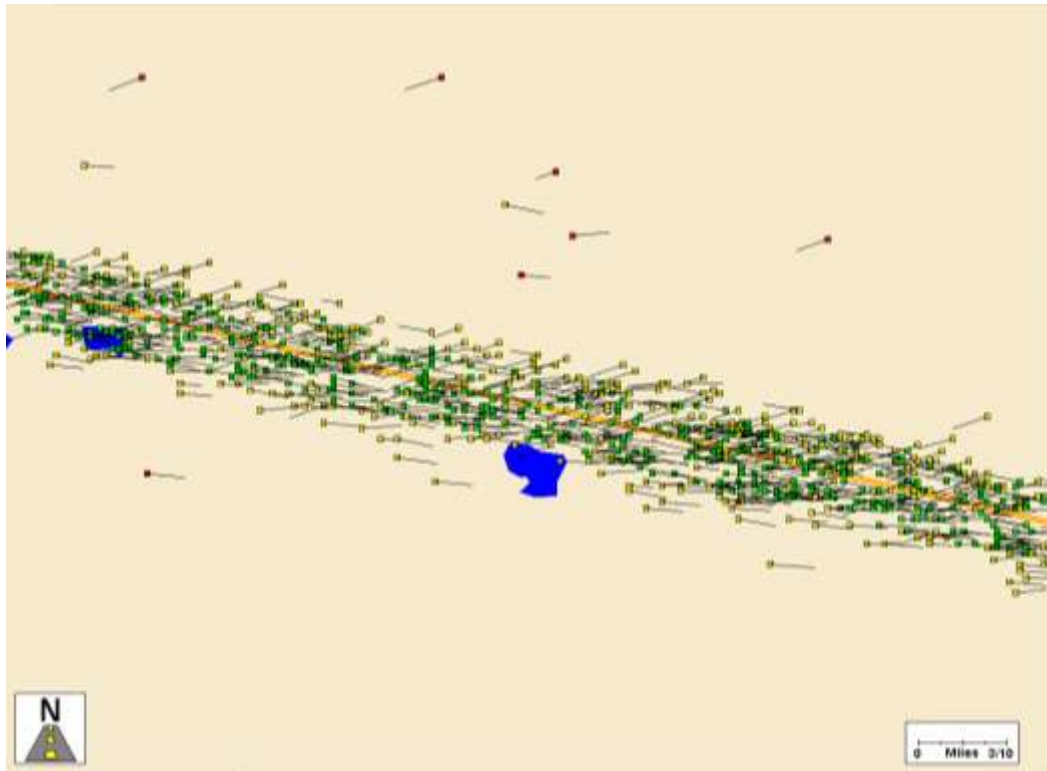

Figure 3, QASPR position data along a stretch of I-80 in Pennsylvania; 13 day sample of over 300,000 trucks.
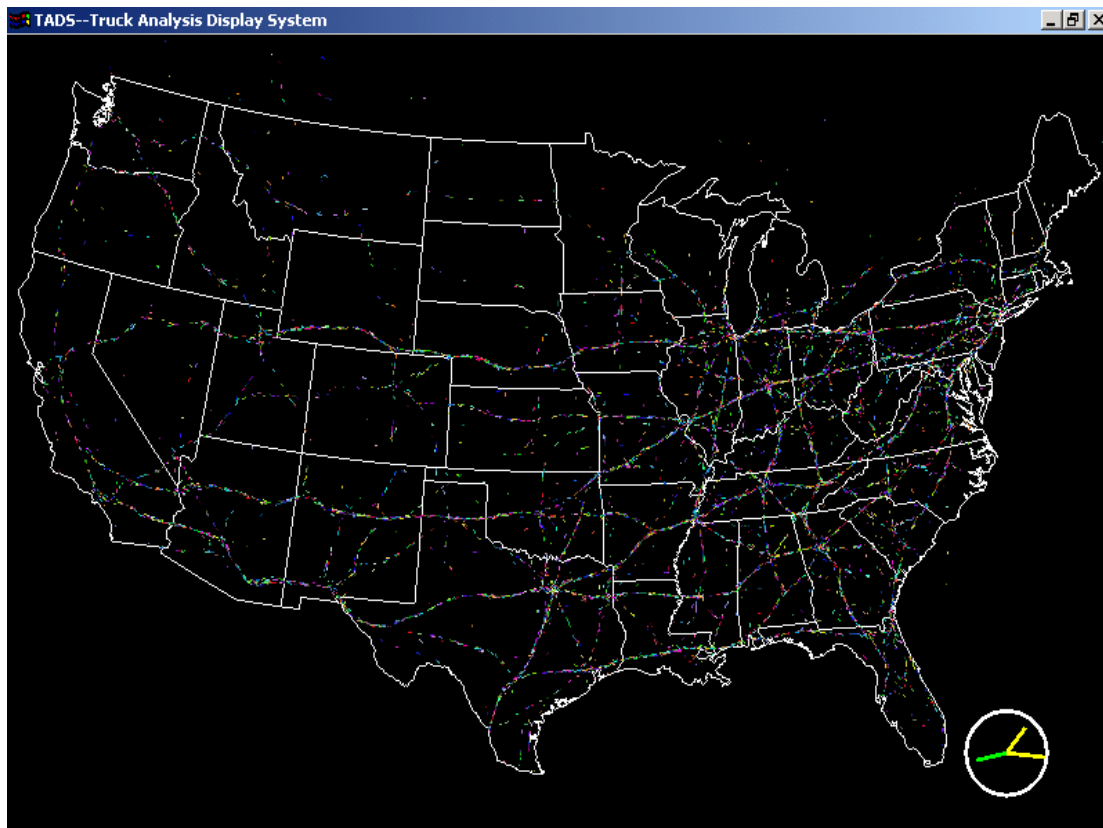
Figure 4, Display of the interpolated instantaneous QASPR location reports of 27,417 trucks using the OmniTRACS system.
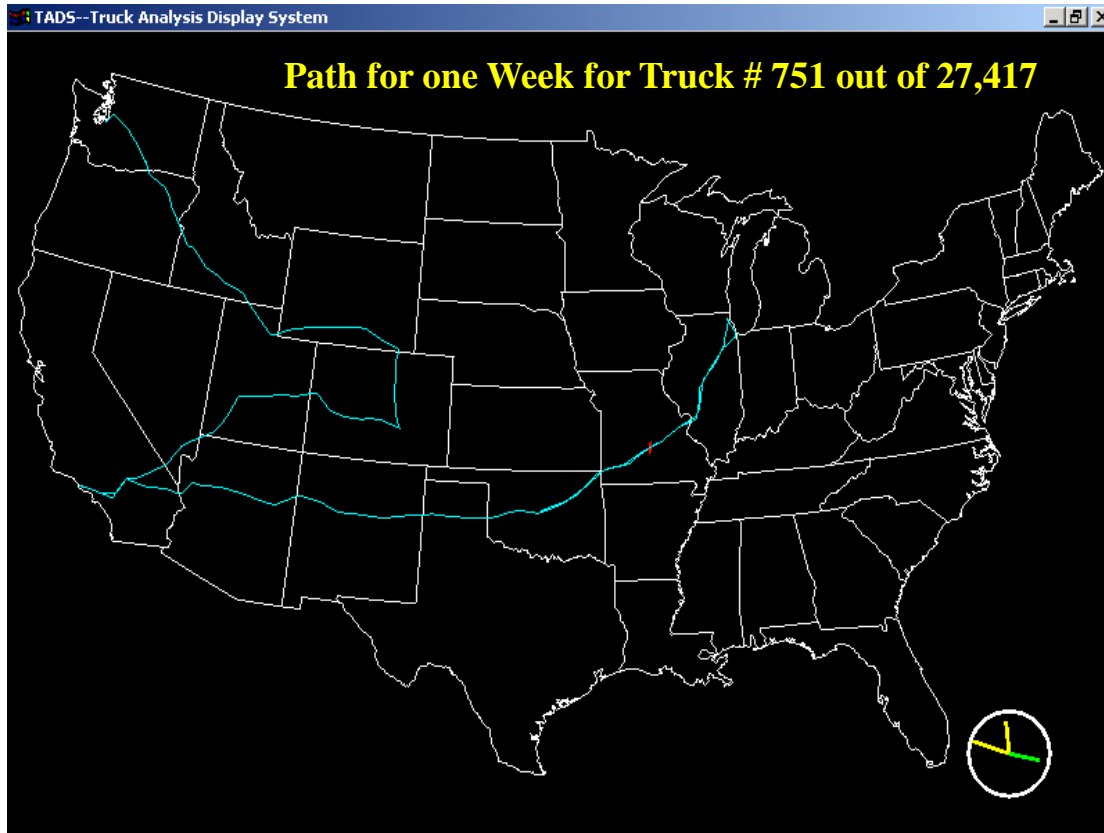

Figure 5, Path for one week of vehicle probe #751.

3.2 WiFi Positioning System (WPS).

Pioneered by Skyhook, WPS uses precise locations of WiFi access points as the underlying reference frame to determine location. http://en.wikipedia.org/wiki/Skyhook_Wireless If sufficient WiFi access points exist nearby, the system is able to determine a position t an accuracy comparable to GPS. These locations tend to be indoors or in dense urban canyons because that's where most WiFi access points have been located and are places where GPS is most challenged. At present WPS can, at best, complement GPS in some dense urban areas. Skyhook's WPS implementation is in the current iPhone.

3.3 Global Positioning System (GPS).

GPS has emerged as the most widely used means of determining position. It provide an inexpensive means of obtaining a rather precise measurement of position and velocity in locations in which the GPS device has an unobstructed view of at least four (4) of the GPS satellites. At rest, accuracy of inexpensive consumer-oriented GPS receivers is nominally stated to be within 15 meters for 3 standard deviations of recordings. This accuracy can be improved to 3 meters using differential correction as provided by the two WAAS geostationary satellites, http://gpsinformation.net/exe/waas.html. Accuracy is also maintained when moving at constant velocity. Small drift errors generally less than 10 meters tend to be introduced when accelerating linearly and turning. Accuracy also diminishes when obstructions such as buildings and foliage cause less than four GPS satellites to be in direct view of the GPS receiver. Such degradation of performance is generally termed the "urban canyon" effect. Since the GPS constellation of 24 active satellites and 7 spares are in 6 inclined circular orbits with 12 hour periods, different satellites move across the sky in different locations at different times. Thus, the urban canyon effect can be inconsistent. In general, the best surface environments for GPS accuracy are the open sea and the open plain. The harshest environments include foliage canopied roads, mountain passes and urban areas such as Manhattan.

In areas free from the urban canyon effect, GPS provides very accurate position, speed and heading data that clearly describe probe vehicle movement patterns. Figure 6 displays the postion data for a FedEx contract carrier over a one week period in 2005. Data reporting rates are nominally every three (3) seconds. Traveled route are readily identified and stop loactions and durations are also readily identified by long sequences of very low reporting speeds. Essentially no GPS drift or outliers are evident in this data set. Problematic locations tend to be cosistently problematic over many vehicles. Problem

areas tend to be those with heavy fioilage or in locations with long overpasses. Tunnels are readily identified because of the conspicuous lack of GPS data over the length of the tunnel and the unusually large scatter of GPS values at the portals of the tunnel. This is caused by the inevitable transition from having 4 or more satellites in clear view on the approachs to the tunnel and having zero satellites in clear view inside the tunnel
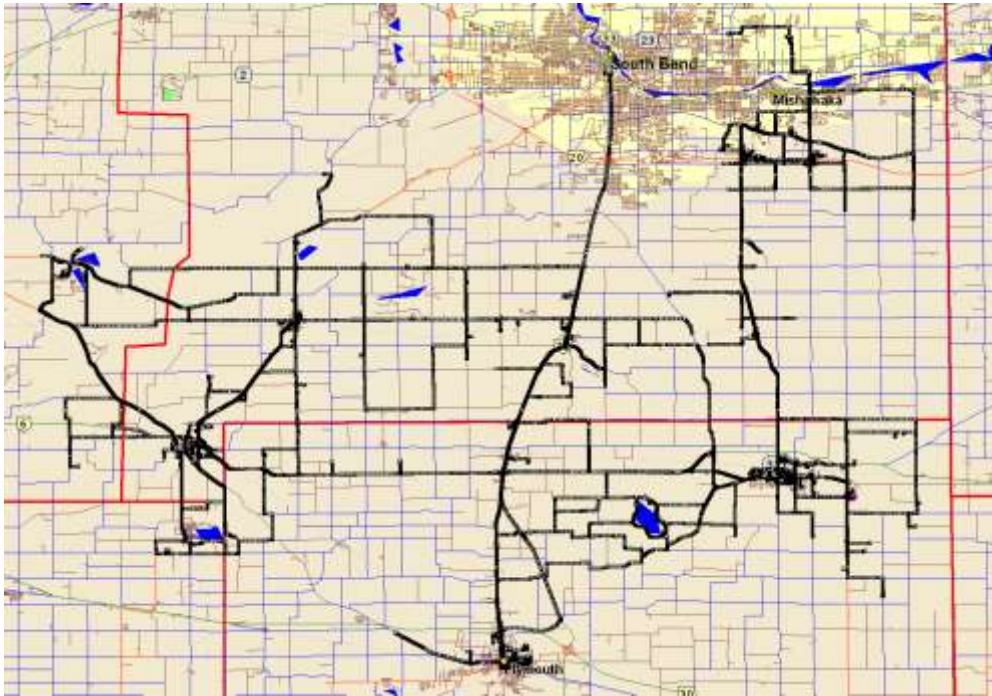

Figure 6, FedEx contract carrier GPS reports every 3 seconds over a one-week period in northern Indiana.

GPS data are captured by a data recording device that contains a GPS antenna and receiver. The antenna receives radio signals broadcast by the constellation of currently thirty-one (31) GPS satellites. The receiver interprets the captures signals, computes a value of position (typically, longitude, latitude and altitude) in a specific reference frame (typically WGS-84), and records the position, time (UTC), date and other data including instantaneous speed and heading. Unfortunately, GPS signals are weak; therefore, an unobstructed line of sight (termed "in-view") is needed to each of four GPS satellites to permit the computation of position in three dimensions. So called 2-dimensional position values obtainable when less than four satellites are in-view are not sufficiently precise to be useful. 3D positional accuracy of even consumer grade GPS receivers have RMS (root-mean-square) values of about 5 meters in latitude and longitude and about 20 meters in altitude even while operating a speeds of greater than 100 kph (kilometers per hour). If the signals are obstructed by entities such as buildings, overpasses or tree canopies, the positional accuracy can substantially degrade. In tunnels, GPS does not function.

Figures 7 and 8 are displays of typical GPS data near the intersection of I-95 and the Mass. Turnpike (I-90) west of Boston at two zoom levels. As can be seen, an overwhelming majority of the position data superimposes tightly on the physical roadway system and the whiskers, which in this case depict the actual speed and heading as reported by the GPS sensor for the corresponding location, are for the most part, as they should be, tangent to the traveled roadway. However, one can definitely see that since the temporary loss of line-of-sight by the vehicles passing under I-90 on the west bound ramp in Figure 8 causes the subsequent positions to shift to the north and have a larger error. Also, the velocity whiskers aren't as tangent to the road surface. Similarly, the data in the opposite direction is scattered to the east after it emerges from passing under I-90. Note also the large concentration of data on the major roadways and ramps and the lack of data on the local roads. This disparity highlights the vast difference in traffic volumes on local roads. This indicates that there may well exist sufficient GPS data to adequately monitor major roads, but there is insufficient data to monitor many, if not most local roads. Luckily, the lack of data on local roads may itself be significant that these roads don't need any active monitoring because they are rarely used , essentially never congested and thus can safely be assumed to always operate at a nominal level-of-service.
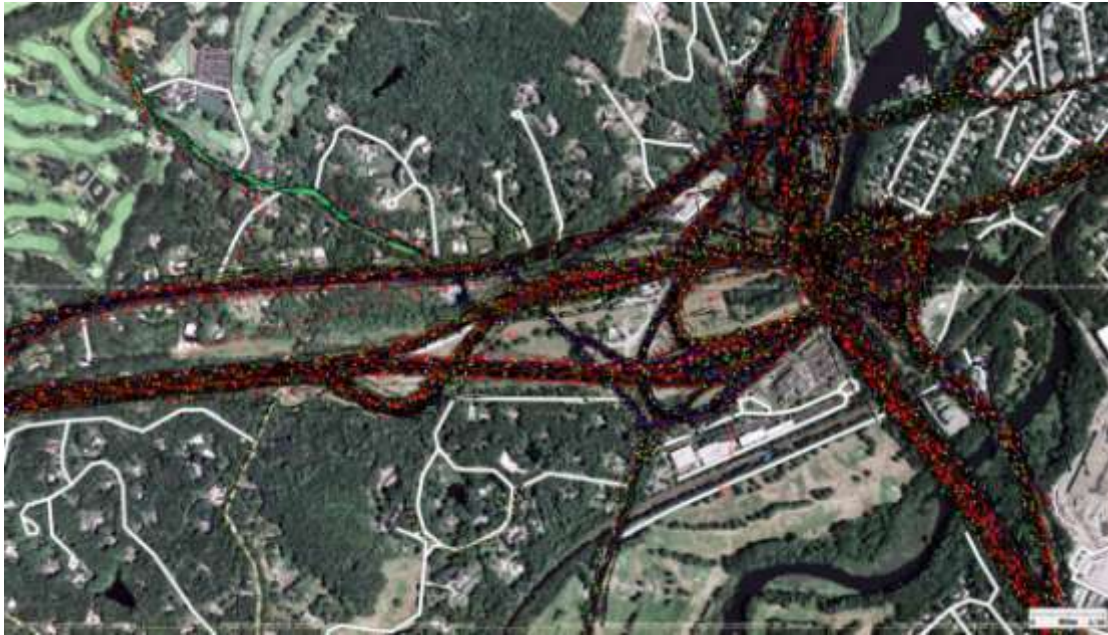
Figure 7, Typical GPS position and velocity data at a major highway intersection, I-95 and I-90 west of Boston
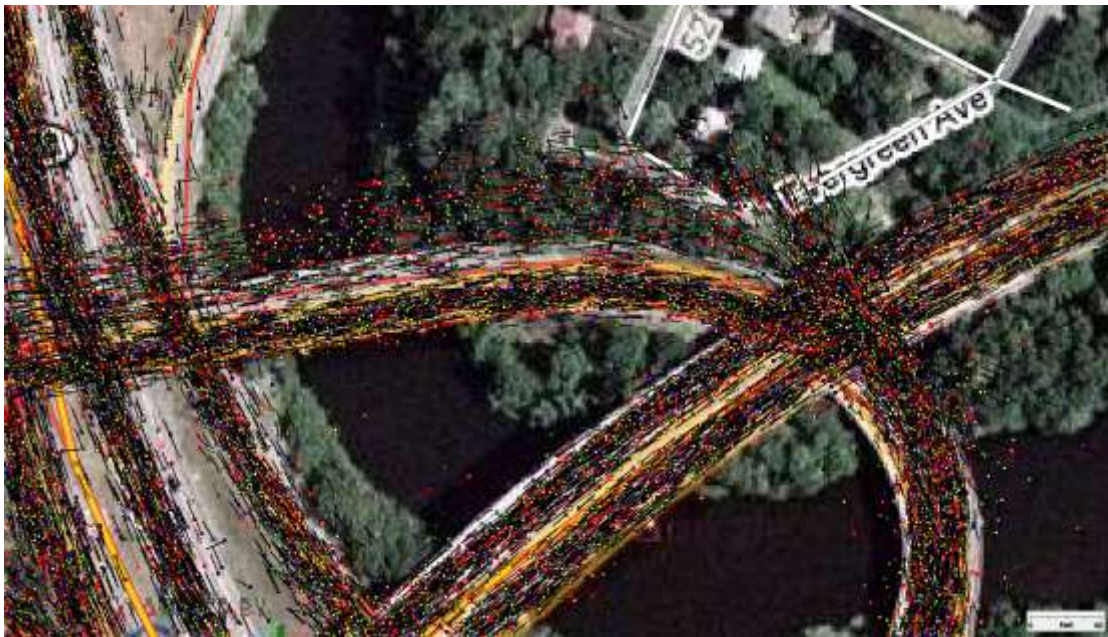
Figure 8, Zoomed-in view of two ramps as they pass under I-90.

In urban areas, the positional accuracy of GPS data can degrade substantially. Figure 9 displays GPS data from vehicles traveling in and around Manhattan. The red points, drawn on top of the black bordered green points are those points whose GPS location differed greatly from its actual position. While there are relatively few of these points in this "worse case", there are a substantial number. They tend to be generated in the midtown area (where tall buildings provide only a narrow access to the sky), the portals of the tunnels, covered roadways (as where the FDR passes under the Gracie Mansion) and bridges (where the steel suspension structures interfere with line-of-sight).

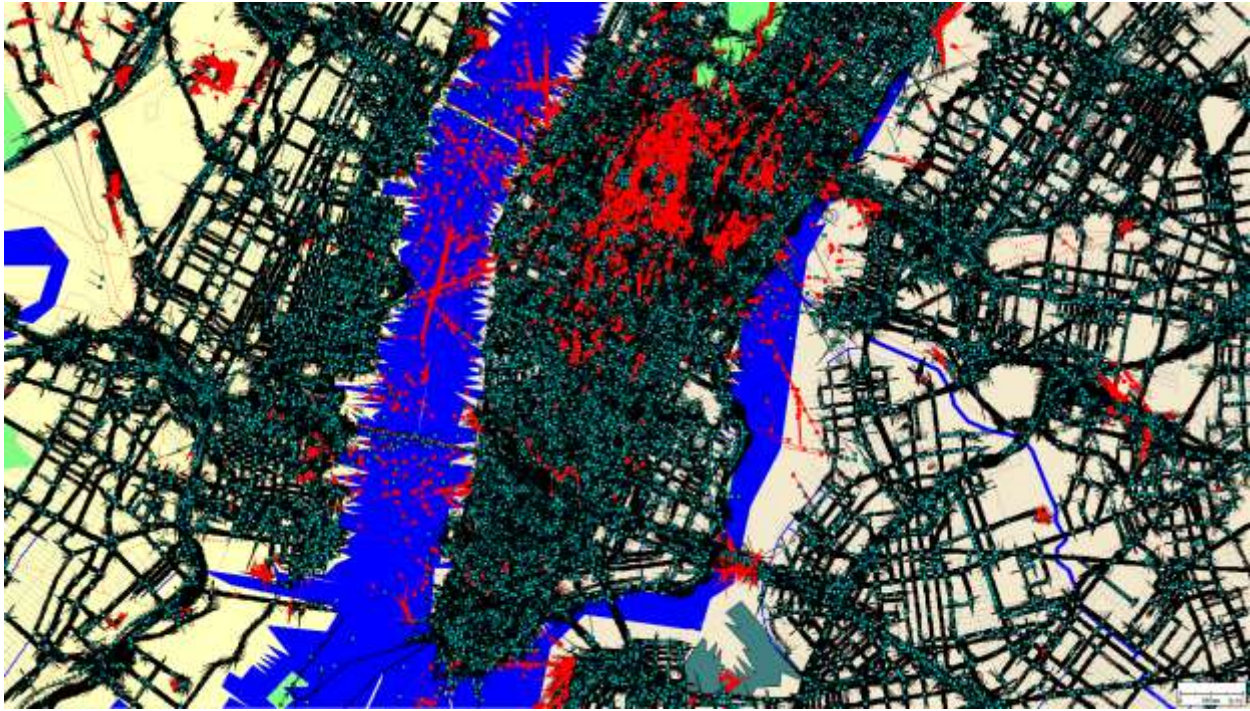**Figure 9, Map-matched GPS positions near Manhattan; red indicates very poor map-match value.**

4.0 Map-matching and measures of travel time

If vehicle probe data, be it QASPR,WPS or GPS, are to be used as a source for monitoring the performance of a roadway segment, the location data corresponding to when the probes were actually traversing that segment need to be isolated. A convenient way to identify any roadway segment is in terms of the sequence of arcs of the digital map database that combine to make up that segment. Thus if the vehicle probe data are tagged with the nomenclature of the arc that the vehicle was most likely traversing at the time of the location record, then all of the data pertaining to any segment can be isolated and analyzed. The process of assigning a most likely digital map location to each vehicle probe data is called map-matching. Bernstein and Kornhauser (1996, 1998) and White, et al (2000) pioneered the use of position, velocity and connectivity with sequentially map-matched locations to form goodness of fit measures for the map-matching process. Quddus' PhD dissertation (2006) provides an excellent summary the various approaches to map-matching and proposes his own version. He made a nominal comparison of the various approaches using a ad hoc data set; however, any comparison must also be reflective of the end use of the map-matched data and incorporate the appropriate formulae that combine the various elements. Moreover, there is no standard probe data set that includes both position data and certified map-matched location. As is often the case, in the absence of such a standard data set, one can't simply do regressions to obtain the best goodness-f-fit function . Instead, experience and art play an important role in defining the best formula. Given that the quality of the map-matching algorithm substantially affects the performance of turn-by-turn navigation systems such as CoPilot, it is not surprising that any subtleties in a map matching algorithm would be kept secret much as is done with the formula for Coca Cola. What can be said is that the map matching algorithm contains an efficient process for selecting the candidate set of arcs to be tested as well as finely tuned coefficients for each term of the goodness-of-fit function . The candidate set is robust in that it includes long arcs while not being burdened with too many arcs. It also efficiently analyzes every shape-point segment of every arc to yield a most likely position along the most likely link in the most likely direction, {Link#, t, dir}, for each and every position data element. Figure 10 shows the map-matched data appended to the position data that has been exposed using a mouse-over command on one of the location data points displayed in Figure 8.
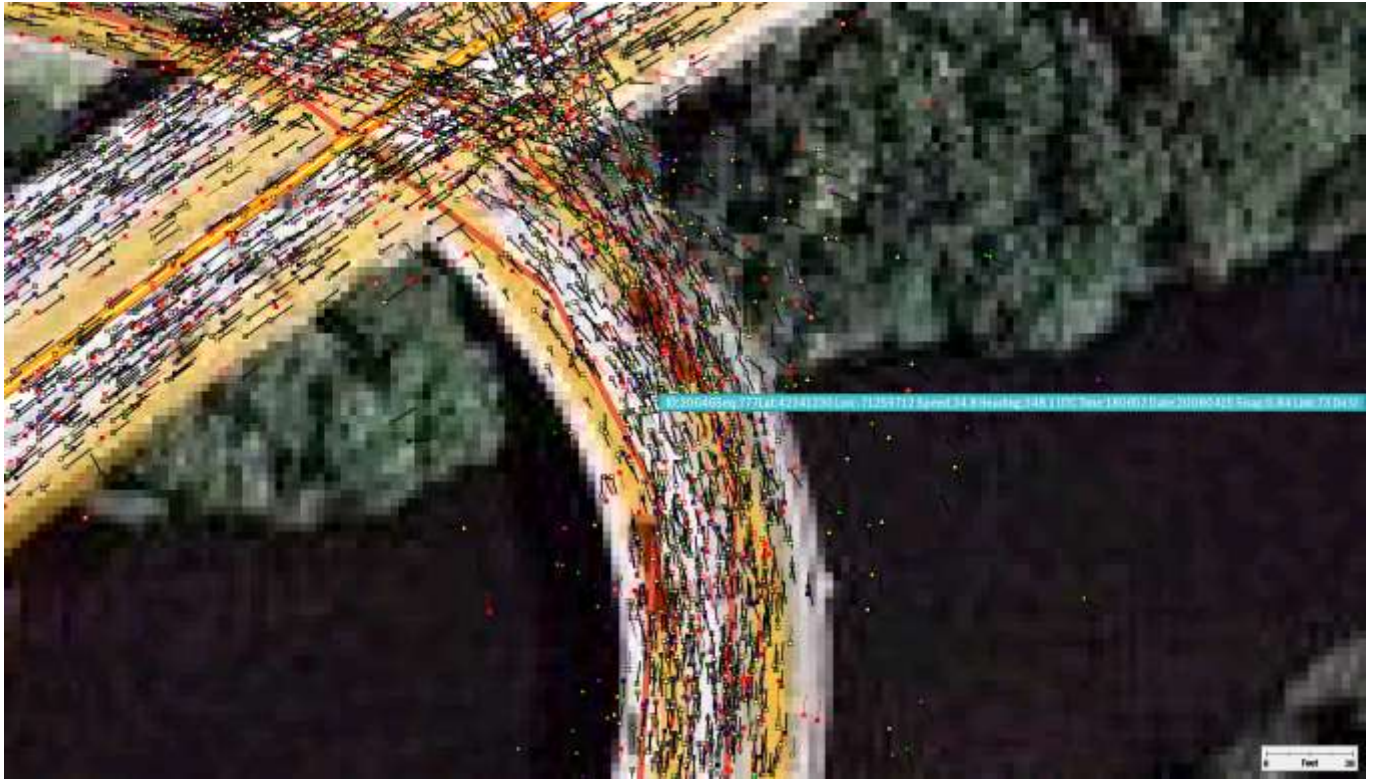
Figure 10, "Mouse-over" exposure example of the nominal and map-matched data assigned to each and every probe data point. Different colors are for different VehicleID.

4.1 Computing Travel Time from map-matched vehicle probe data.

The map matching process places a vehicleID at a specific digital map location at a specific date and time. The travel time experienced by any vehicleID between any of its map-matched locations is readily computable. Unfortunately, the locations along any segment are essentially random and not organized to be at specific locations. However, by restricting the computation of travel time to map-matched position reportings that are nearest to specified locations, the travel time between these points can be considered to be representative of the travel times between the specified locations. If the vehicle speed can be assumed to be constant in the neighborhood of a specific point, then the time of passage by any other point in that neighborhood is obtained by a simple linear extrapolation from the specific point using the constant speed. Usually, speeds tend to be more constant at the middle of arcs that span branch points in a digital map database. Consequently, if one has the freedom of choosing the start and end of a highway monitoring segment, the preferred choice would be to choose midpoints of arcs ($t_k = 0.5$) rather than any other specific point. These preferred locations are called "Monuments". Also, the monument arcs need to be long enough and the probe data rate large enough such that there is a high probability that the Monument arc will have multiple map-matched position records. For example if the data rate is once every three seconds, then any freeway links longer than $1/10^{th}$ of a mile (500 feet or 200 meters) can be expected to have at least two map-matched points per link. Local streets as short as $1/20^{th}$ of a mile (250 feet or 100 meters) can be expected to have at least two map-matched points per traversal. This approach allows for a substantial reduction in data points by the elimination of all but the map matched point that is closest to the midpoint during any traversal of the Monument arc. This reduced dataset is called the OneMon dataset. Figure 11 shows only the OneMon elements at the intersection of I-8 and I-895 near San Diego. Very short arcs make up this section of ALK's digital map database, thus the data reduction is not much greater than 50%. On the other hand, in more rural areas, such as along CA 53 in southeast California, shown in Figure 12, the reduction is a factor of 10.

Figure 11, Display of only the most centered GPS point for each vehicle traversal of each arc of ALK's digital map database at the intersection of I8 and I 805 near San Diego, so-called "OneMon" data set.

Using the concept of monuments, performance can be measured and monitored for road segments rather than individual points, where the road segment is the physical infrastructure between pairs of monuments, so called "Monument2Monument",M2M , pairs. While monuments could be located anywhere along an arc, unless otherwise stated, monuments are taken to be at the midpoint of an arc. Computation of travel time between any two "Monuments" involves geometry as depicted in Figure 13. M_i is the midpoint location of the upstream arc and M_j is the midpoint location of the downstream arc. If M_i and M_j are adjoining, then the distance between M_i and M_j , $D_{Mi->Mj}$ is simply

$$D_{Mi->Mj} = \frac{1}{2} (D_i + D_j) \text{ where } D_k \text{ is the length of arc k.} \tag{2}$$

Figure 12, Display of only the most centered GPS point for each vehicle traversal of each arc of ALK's digital map database along CA 53 in Southeast California, so-called "OneMon" data set.

If there are intervening arcs, then their length is simply added to equation (2). The distance between a map-match point and the arc midpoint is simply $|t_k - 0.5| * D_k$ for any arc k. The proper addition or subtraction of this values for k = i and k = j yields the distance between the map-matched points, $\underline{D}_{i\to j}$. The measured average speed, $\underline{S}_{i\to j}$, is simply

$$\underline{S}_{i\to j} = \underline{D}_{i\to j} / (\underline{UTC}_j - \underline{UTC}_i) \quad \text{where } \underline{UTC}_k \text{ is the UTC time of the } k^{th} \text{ map-matched point.} \quad (3)$$

The Travel Time from M_i to M_j,

$$TT_{i\to j} = \underline{D}_{i\to j} / \underline{S}_{i\to j} \quad (4)$$

Computation of DelDistance

M — Monument ; located at midpoint of link, Length of link is Pre-computed is distance Mi to Mj

■ — Map matched location of GPS point
Known is link # and % distance along link

★ — GPS point location

Mj

Mi

Distance between ■ and **M** is computed properly added or subtracted from distance to yield DelDistance
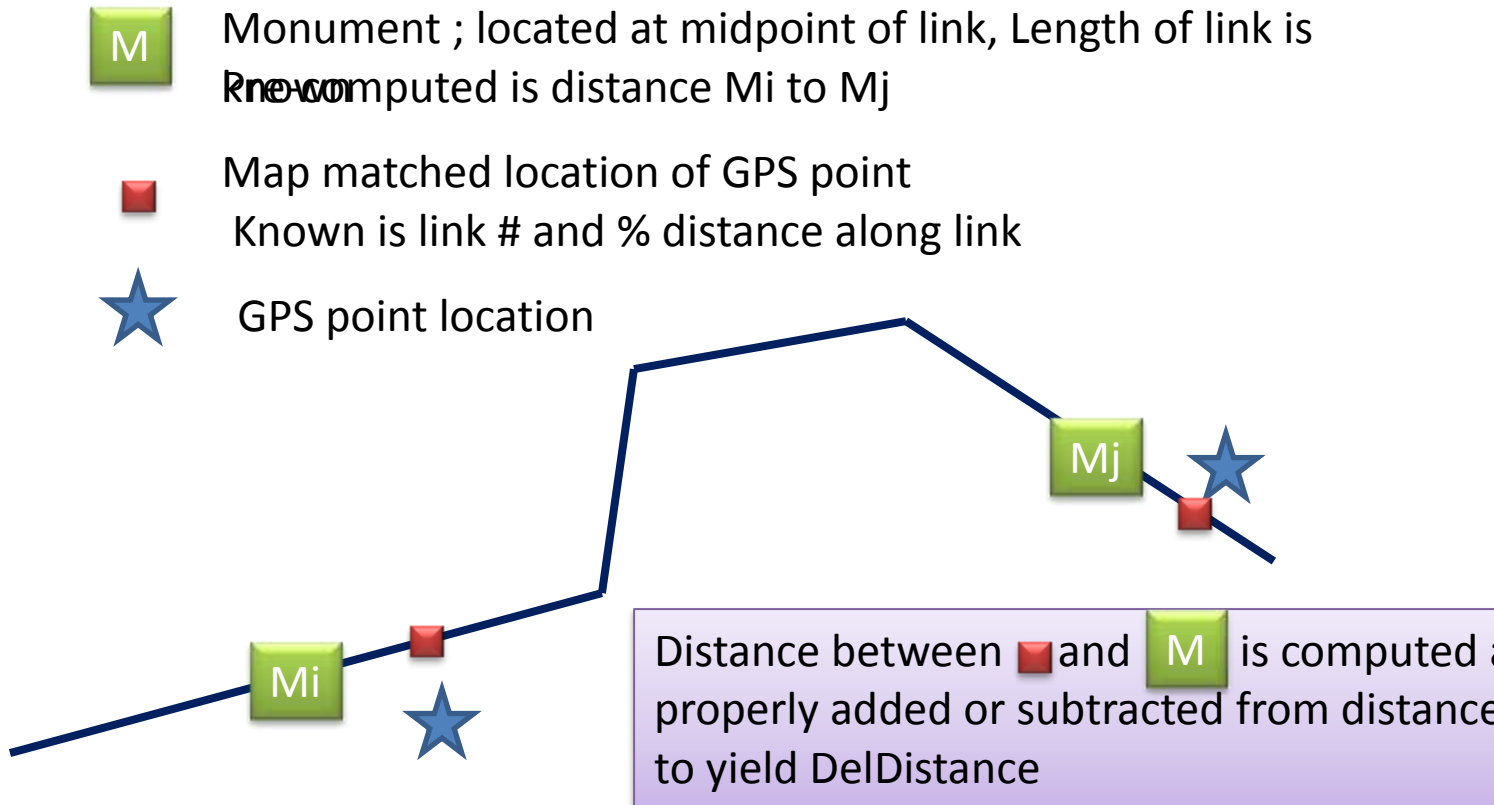
Figure 13, Display of data geometry of typical M2M pair.

4.2 Measures of travel time and speed for M2M pairs.

Given any M2M pair it can be expected that more than one observation of travel time and measured average speed will be available for any classification of the data. If large amounts of data exist, then the data may be classified according to time-of-day and possibly even day-of-week. In any event, the data for any classification will be characterized by the following performance measures:

1. its cumulative probability distribution which is obtained by plotting the classified data in sorted order,
2. its median value (in order to avoid the biases caused by outliers) and
3. Its standard deviations.

5.0 Example using GPS probe vehicle position data to monitor the speed performance throughout the Windsor-Québec corridor

Probe vehicle GPS data formed the basis of a performance monitoring demonstration for the 1,127.7 km Windsor – Quebec corridor; see Figure 14; Kornhauser and Batchu (2009). Time-sequenced GPS location data, termed "GPS tracks", logged during each trip of a fleet of 4,950 unique probe vehicles during the first half of 2008 formed the basis of the corridor monitoring, as summarized in Table 1. Each of the 60.66 million GPS data points, typically 2 minutes apart on any GPS track, contained NMEA position (latitude and longitude in WGS-84 CS), date and UTC. Instantaneous speed and heading were not available. An average speed (GPSSpeed) and heading to next point were computed.

# unique Device_IDs	Total GPS data points	Total travel hours	Total distance traveled (km)	Average distance traveled (km/ Device_ID)	Average travel time (hr./ Device_ID)
4,950	60,659,746	1,345,475	118,357,762	23,910	271.8

Table 1, Summary of GPS data for unique Device_IDs (trucks) for 6 month data sample.
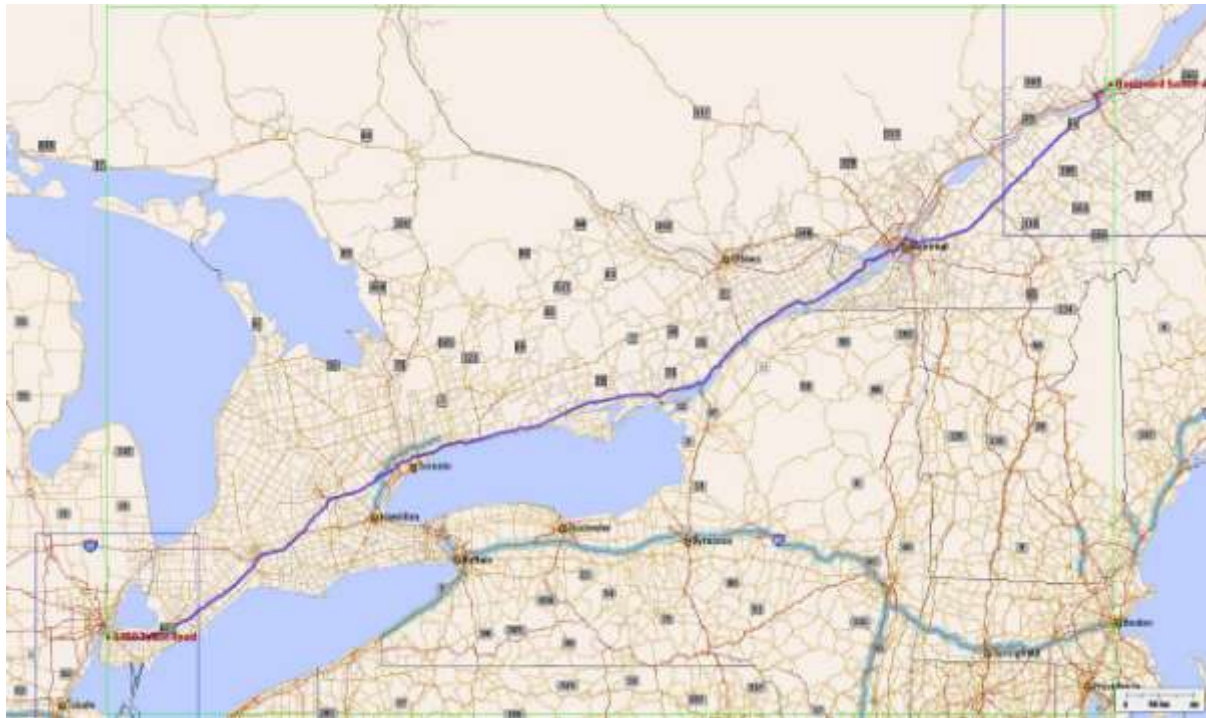

Figure 14 Map of Québec-Montréal-Windsor Corridor; 1,127.7km.

Figure 15 and 16 display geographically the GPS data for January 2008 which are typical of those that exist for each of the other five (5) months. Figure is a broad geographic overview of the GPS point locations. Location data are seen to exists throughout the eastern portion of the Canadian National Highway System. At this level of geographic precision, numerous data points are super positioned. Each is color-coded by and drawn in order of Device_ID and Trip _ID. Trip_IDs were derived by a date and time sort of all of the data for each unique Device_ID. When traveling, the data for any Device_ID tended to be separated by about 2 minutes. Data points with speeds less than 5 kilometers per hour were originally purged from the dataset and not included in the counts contained in Table 1. Gaps longer than 20 minutes were interpreted to mean that the vehicle had stopped travelling for what could have been any of a multitude of reasons having nothing to do with the fundamental performance of the roadway system; for example, making a pickup or delivery, getting fuel, and/or taking a break, but not involved in a stoppage of traffic. This tacitly assumes that at no time during these six months did traffic come to rest for greater than 20 minutes on any roadway segment. These gaps were used to break up the data for each Device_ID into a series of Trip_IDs. For each Device_ID, the first GPS point as well as well as the first GPS point after a gap of greater than 20 minutes were taken as the beginning of new Trip_ID and the last GPS point prior to any gap greater than 20 minutes and the overall last GPS point were taken as the end of a Trip_ID. In this way a unique Trip_ID was assigned to each GPS point for each Device_ID. When drawn, only the color of the last drawn Trip_ID for the last Device_ID at any point survives. By inspecting continuations of identical colors, one can see that individual Trip _IDs traversed substantial distances.

Figure 16 is a close-up view of the January 2008 GPS data near the interchange of the MacDonald-Cartier Freeway (ON 401) and ON 400. Displayed are the locations of individual Device_IDs (latitude, longitude) as well as a black line (a "whisker") whose length is proportional to the "reconstructed using the next GPS location" GPSSpeed. Direction is towards the location of the next GPS location for this Trip _ID. As can be seen, numerous data points exist in this location. Since the data rate for any Trip _ID is roughly every two minutes, any trip is unlikely to have more than one data point in this figure; thus, one can readily see the large number of independent observations that are available for monitoring the performance in this part of the corridor. Figure 17is a geographic bandwith chart of the number of unique Trip _IDs by segment along the Windsor-Montréal corridor. The width of the green bands along the corridor is proportional to the unique number of Trip_IDs providing GPS data on those segments of the corridor. The bandwidths have been drawn using a right hand rule. When facing in the direction of the traffic, the width of the band to the right of the black centerline of the roadway segment is proportional to the volume in that direction. Note that, as expected, data rates are similar in each direction. The data are especially concentrated throughout the corridor from Québec to/from Windsor relative to the other areas of Eastern Canada. It is extremely heavy in the Montréal-Windors portion of the corridor which tend to have about 4,000 unique Trip _IDs per direction per month and is heaviest in the Toronto-Hamilton area where the rate jumps to about 10,000 per direction per

month. It is substantially less intense in the Québec-Montréal corridor, having fewer than 1,000 unique Trip_IDs per direction per month.

Figure 15 Geographic display of GPS data from January 2008 on the Québec-Windsor corridor and surrounding major roadways (many points a super positioned)

Figure 16, Geographic display of GPS data including a heading vector whose length is proportional to the speed assigned to the point. Data for Jan. 2008 at the intersection of the MacDonald-Cartier Freeway (ON 401) and ON 40.
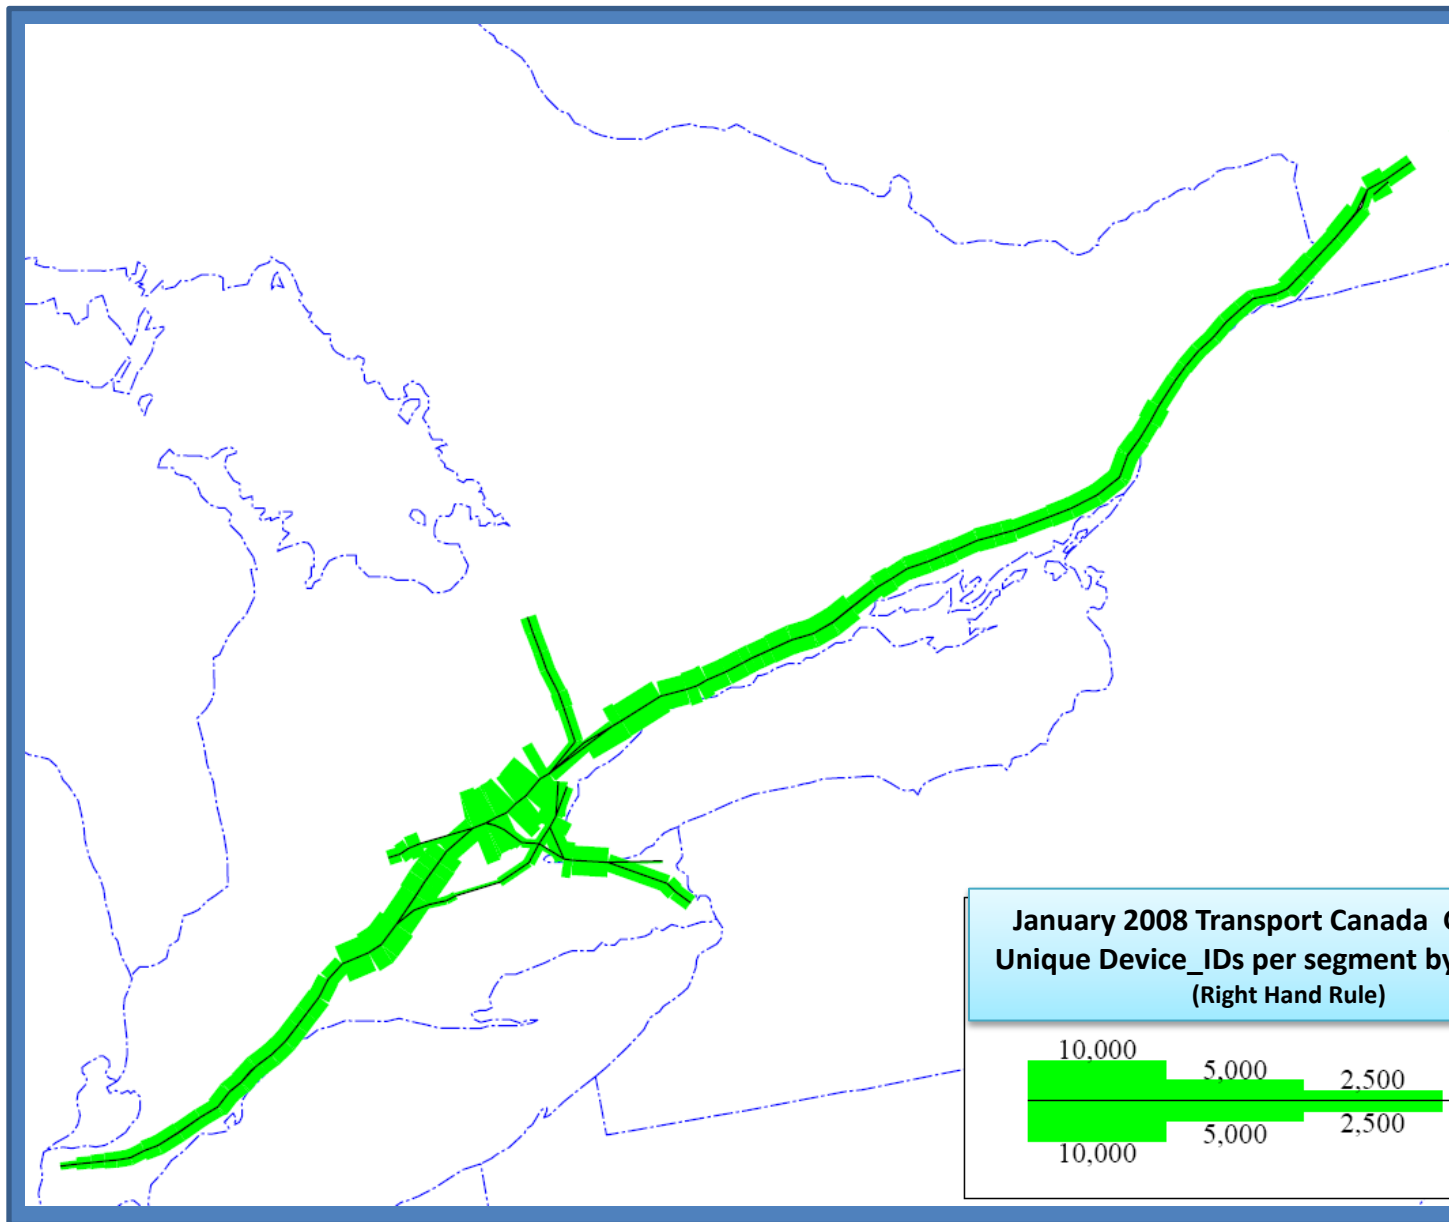
Figure 17, Geographic bandwidth chart of the number of unique Device_IDs by segment along the Windsor-Montréal corridor.

Figures 18 and 19 show the GPS point of two Device_IDs during the first half of 2008. Figure 18 is for Device_ID 9992. Its travels are mainly up and down much of the Montréal-Windsor (M-W) corridor. Figure 19 is for Device_ID 9999 whose travels are even more constrained to the M-W corridor.

Figure 18, Travels of Device_ID 9992.

Figure 19, Travels of Device_ID 9999.

5.1 Selection Québec-Windsor Monuments (QWM) and Monument-to-Monument (M2M) monitored segments.

Speed on the Québec-Windsor corridor was monitored across Momument-to-Monument (M2M) segments that combined to span the 1,127 km corriror. Each M2M sgment was defined as stretching form the midpoint of acrs in ALK's "Level l (excludes interchanges with roads used mostly for local traffic) " digital map representation of the corridor whose length was

at least 3.5 km. This implied that even when traveling at the highest speed, there was a very high probability that any vehicle traversing the monumnet link would have at least one GPS position appearing in the database. These arcs are called QWM links. Figure 20 is a display of one of the QWM links, # 851. Figure 21 displays the "OneMon" GPS point closest to the center of that link for each TripID that traversed that link. So many data points are superpositioned on this short Monument link, the actual decrease in the concentration of data points at the endsof the link it is not obvious .

Figure 20, Display of QWM link number 851 (43.601144N, 79.777827W).

Figure 21, Display of SingleMon GPS data for QWM link number 851 (43.601144N, 79.777827W).

Because the GPS data density was singnificantly different, the Québec-Windsor corridor was seperated into a Québec-Montréal corridor and a Montréal -Windsor corridor. The Québec- Montréal corridor, 269 km in length, has 29 monument links that form 28 contiguous M2M segements. The Montréal -Windsor corridor, 858 km in length, has 91monumnet links that form 90 contiguous M2M segments. During the analysis, it became apparent that many of the segments were on roadway stretches in which there was no congestion. The tedium involved with displaying the detailed analysis of all 90 segments, most of whom showed no speed reductions was counterproductive. Consequently, it was decided to constrain the analysis to 28 M2M segments as was being done for the Québec- Montréal corridor. In selecting the 28 segments, the segment west of Montréal and immediately east of Windsor were selected as well as all of the contiguous M2M segments near Hamilton and Toronto. The remaining segments were distributed somewhat evenly between east of Toronto and west of Hamilton. Figure 22 shows one of the monitored M2M segment, number 851,852.
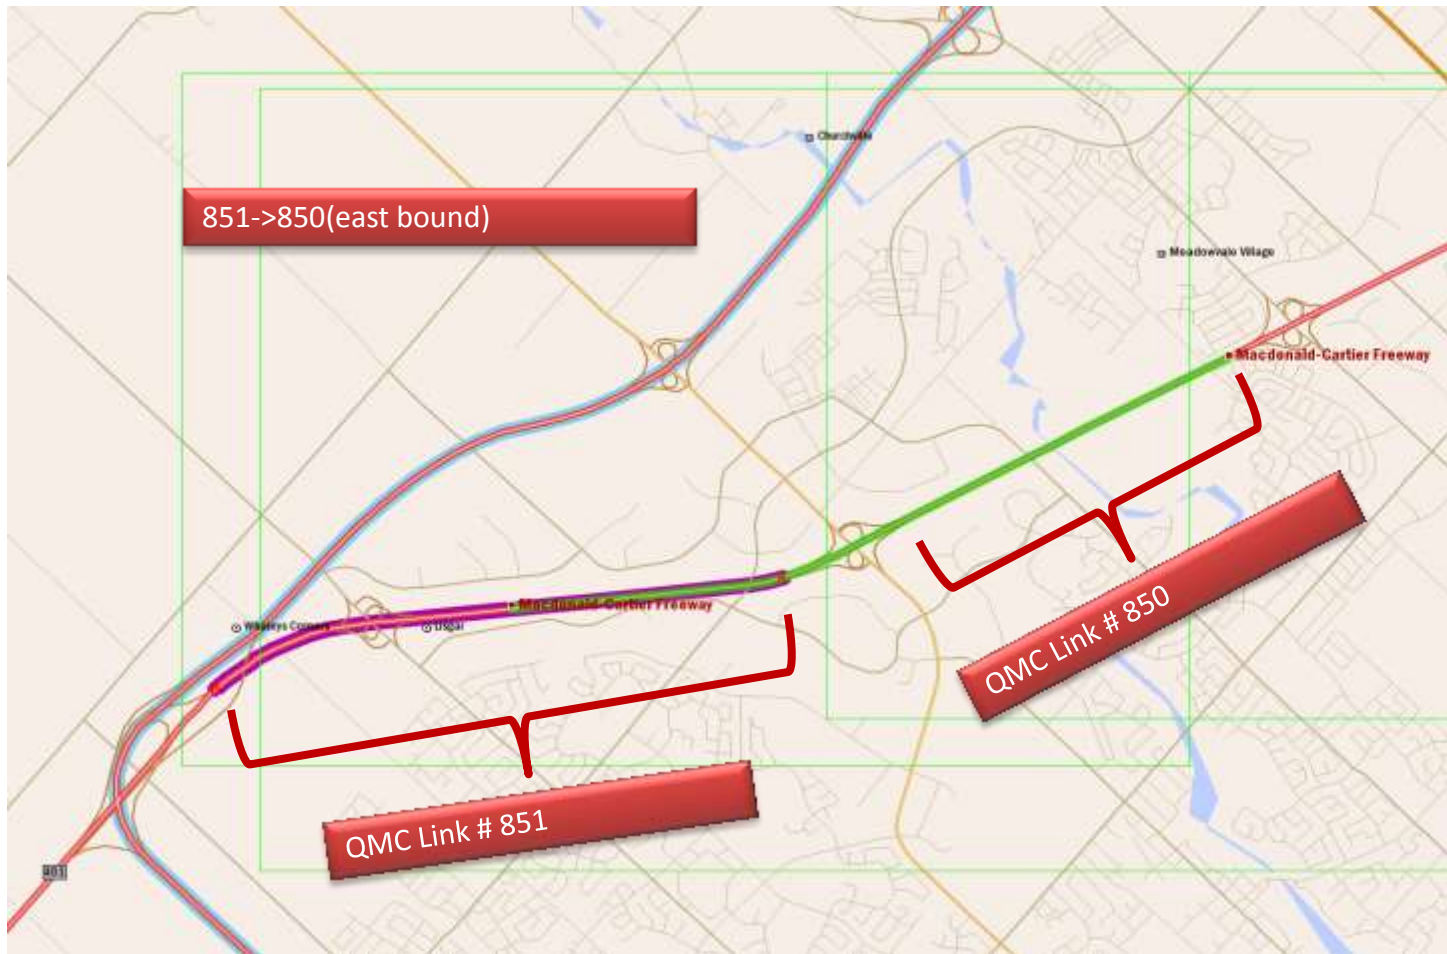
Figure 22, Geographic display of QWM 851 to 850 eastbound, west of Toronto the MacDonald –Cartier Freeway

5.2 Analysis of data rates and measured average speed throughout the Windsor-Québec corridor.

The 60,659,746 GPS data points in the original 6 month data file produced 2,023,184 individual measured average speed (called m2mSpeed) records for the 28 QWM segments that combine to span the Québec-Montréal (Q-M) and the 28 QWM segments that appropriately sample the Montréal-Windsor (M-W) corridor. An "observation" is defined as one of the 2,023,184 individual m2mSpeed records. These observations are not at all uniformly distributed across the corridors. Substantially more observations (87%) are on the M-W corridor than on the Q-M (13%). Distribution of observations per average day across each corridor by travel direction can be seen in Figure 23. Observation rates are low and fairly uniform across the Québec-Montréal corridor. They tail off substantially to unreliable levels as one enters Québec. Apparently few of the trucks providing data for this analysis enter the center of Québec City. There is also a drop in truck traffic at the western side of Montréal. There are several routes through and around Montréal that effectively split and distribute truck traffic to several alternative routes rather than concentrate on one route. This may be the cause of the drop of truck traffic on the western side of Montréal. In the Montréal –Windsor corridor the largest observation rates, reaching values greater than 300 observations per average day, are found in the segments along CN 401 near Kitchener and west of Toronto.

5.2.1 Analysis of data rates.
Observation rates are neither uniform by time-of-day nor day-of-week. Figure 24 shows the distribution by hour-of-day of the number of observations for an average weekday. Given that the average values are substantially greater than 10 for every hour except those in the late evening, many, if not all, of the M-W segments can be expected to be substantially represented on a time-of-day basis during weekdays. This is not true for the Q-M corridor. Here values are consistently less than 5, thus it can be expected that there will not be sufficient observations to properly characterize traffic flow speeds on a daily basis except at a few particularly dense segments at only some hours of the day.

Figure 25 shows the distribution by hour-of-day of the number of observations for an average weekend day (Saturday and Sunday). Credible hourly performance for any weekend day is not justified for the Q-M corridor and may be barely justified for some segments for some hours during some days. What is sufficiently represented is the increase in observations in the later hours on a weekend day. This increase occurs only on Sundays, not Saturdays and is consistent with the increase

in trucking activity late on Sunday in preparation for the pickup and delivery activities on the following Monday morning. Distributions on the 5 holidays during the first half of 2008 were similar to weekend distributions.
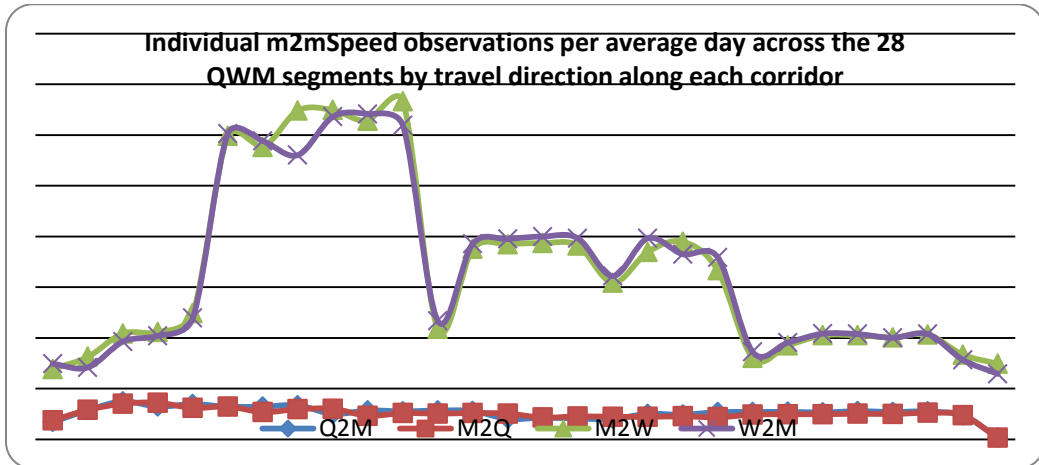

Figure 23, Individual observations per average day across the 28 QWM segments by direction along each corridor.
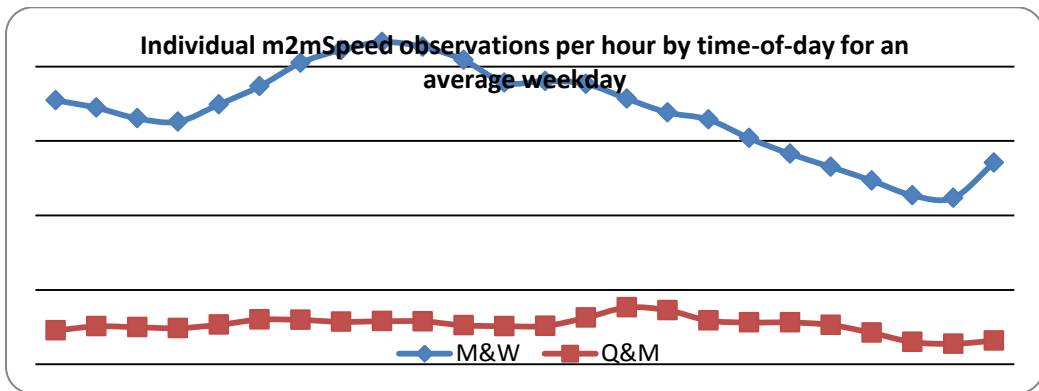

Figure 24, Individual observations per hour by time-of-day for an average work weekday (M-F)
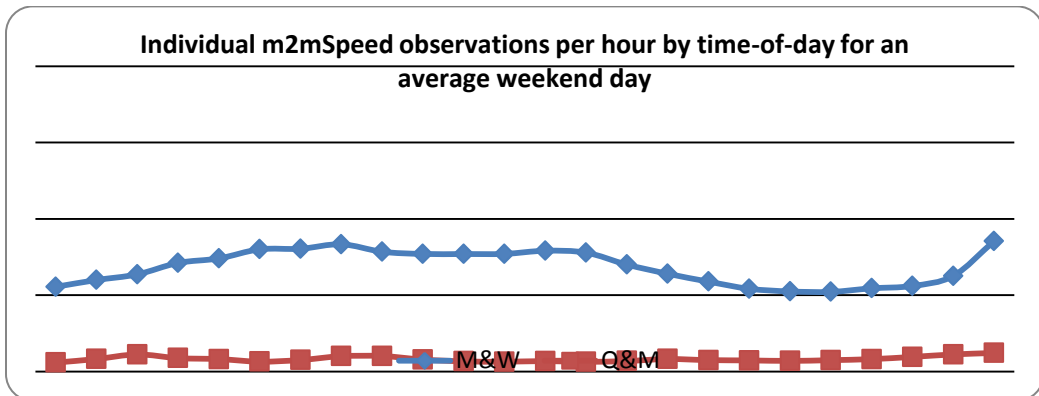

Figure 25, Individual observations per hour by time-of-day for an average weekend day.

5.2.2 Analysis of measured average speed (m2mSpeed). For the M-W corridor, sufficient data existed to analyze m2mSpeeds by hour throughout the six months. Many of the segments showed little diminished speed throughout most days. Figure 26 and 27 show the data for two typical weeks for M2M segment 686 to 687. Figure 26 shows a substantial slowdown for a short period during April 10 (100[th] day since the beginning of 2008) and Figure 27 shows only one brief and not very intense slowdown during the early morning hours of April 14. Otherwise, all of the data are in a narrow band between 90 and 110 kph with median values near 100 kph. The variance can be mostly assigned to individual driver preference rather than congestion or any breakdown in segment performance. This is highlighted by the cumulative

24

distribution of all of the m2mSpeed data points for each of those weeks. This is represented by the monotonically increasing series of red points. One can see a very narrow upper tail, indicating that this gps sample has very few high speed data glitches or fast driving "cowboys". The slightly more prominent lower tail indicates that there were very few truckers caught in any performance slowdown in this corridor. The concentration of the slowdown data in one particular point in time strongly suggests an incident (weather, accident or temporary road repair) was the likely cause of the slowdown. No recurring congestion is observed.
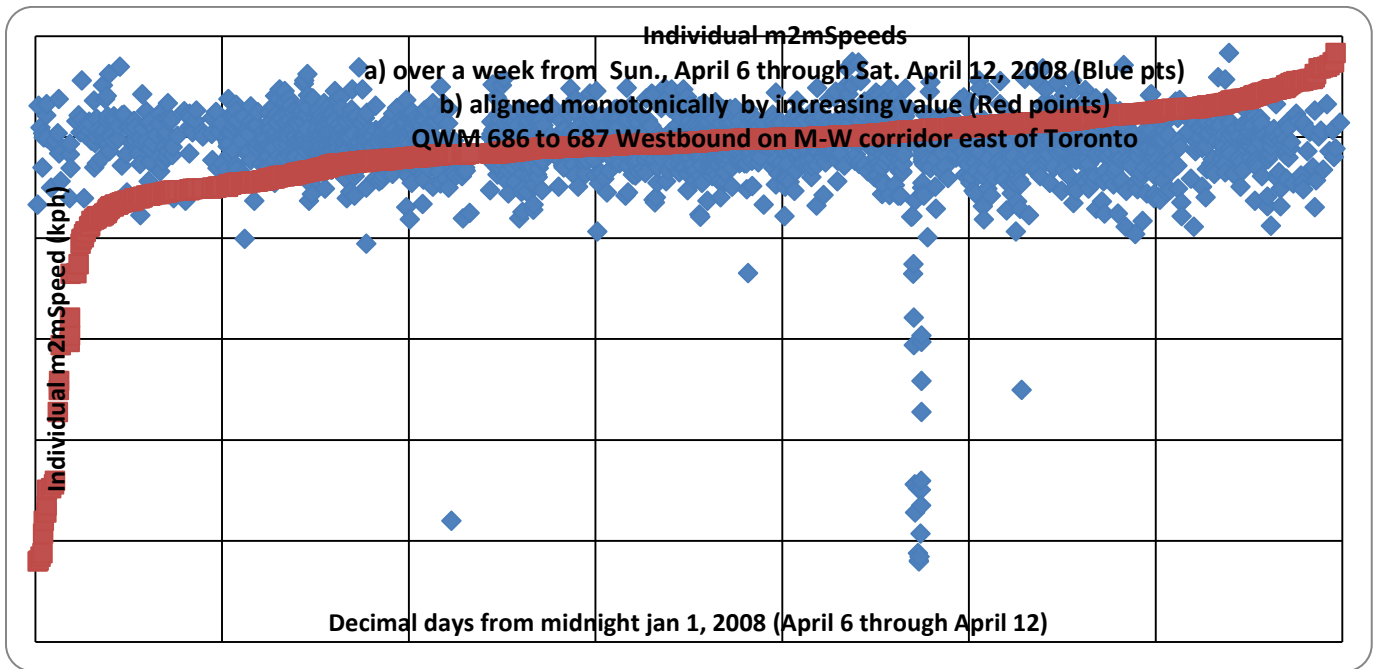

Figure 26, Individual m2mSpeeds distributed by time-of-day and by increasing value for the week beginning April 6, 2008 westbound on QWM 686 to 687.
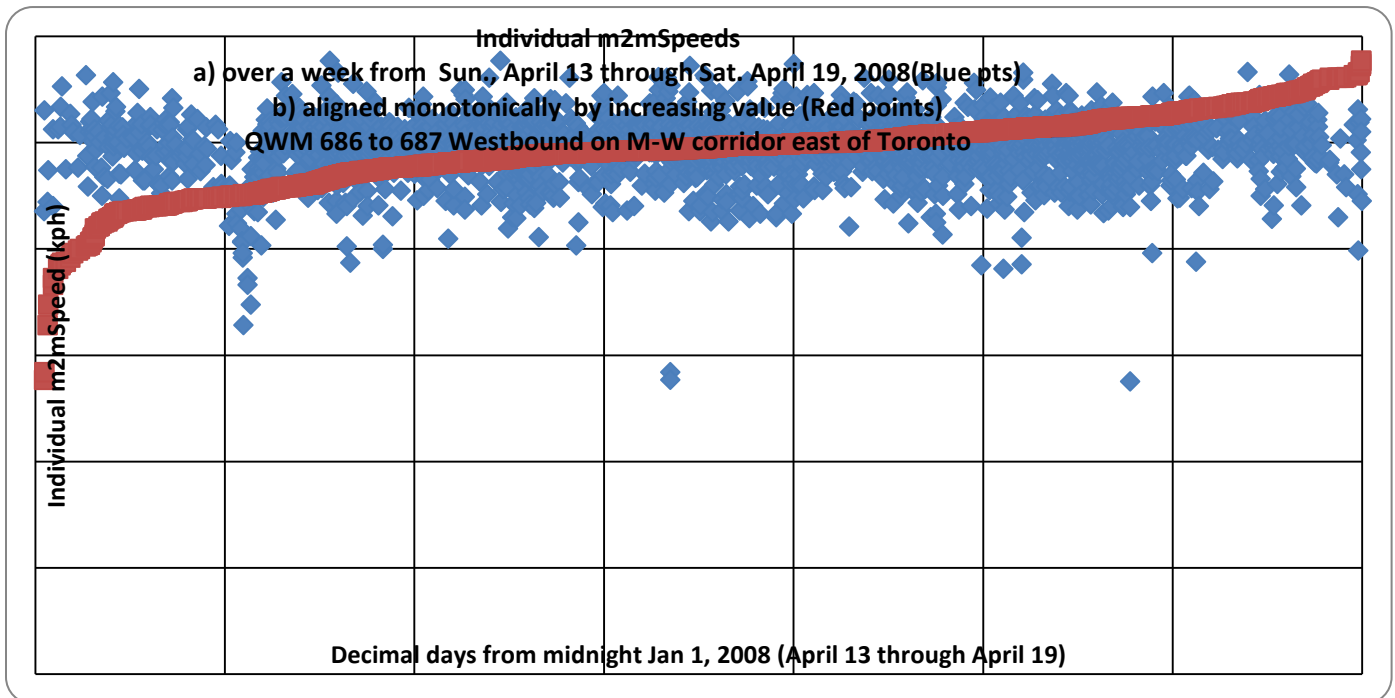

Figure 27, Individual m2mSpeeds distributed by time-of-day and by increasing value for the week beginning April 13, 2008 westbound on QWM 686 to 687.

Substantial recurring speed performance degradation is exhibited on M2M segment 851 to 850. This is shown in Figures 28 and 29 for the same two weeks as Figures 26 and 27. One can readily see the recurring reduction in speed performance on

25

weekday mornings and afternoon. Fridays have the harshest speed reductions in both the morning and afternoon. Monday is also harsh; however, some of that may be due to weather because M2M686, 687 also show a little speed reduction during Monday April 14. Saturday has a hint of midday congestion and Sunday is essentially free of any loss in speed performance. The intensity of the performance reduction on the observations is highlighted by the cumulative distribution of speed (red points). The lower tail, representing almost 40% of the observations experienced congestion delay.
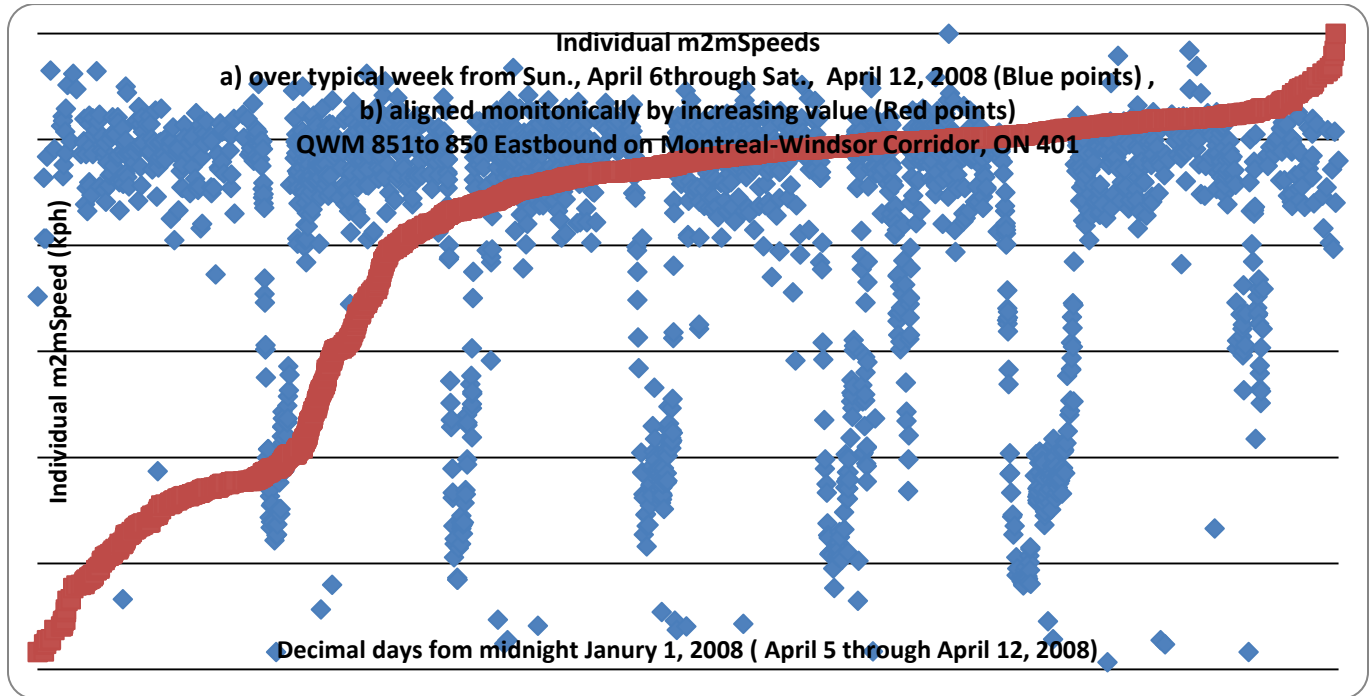

Figure 28, Individual m2mSpeeds distributed by time-of-day and by increasing value for the week beginning April 6, 2008 eastbound on QWM 851 to 850.
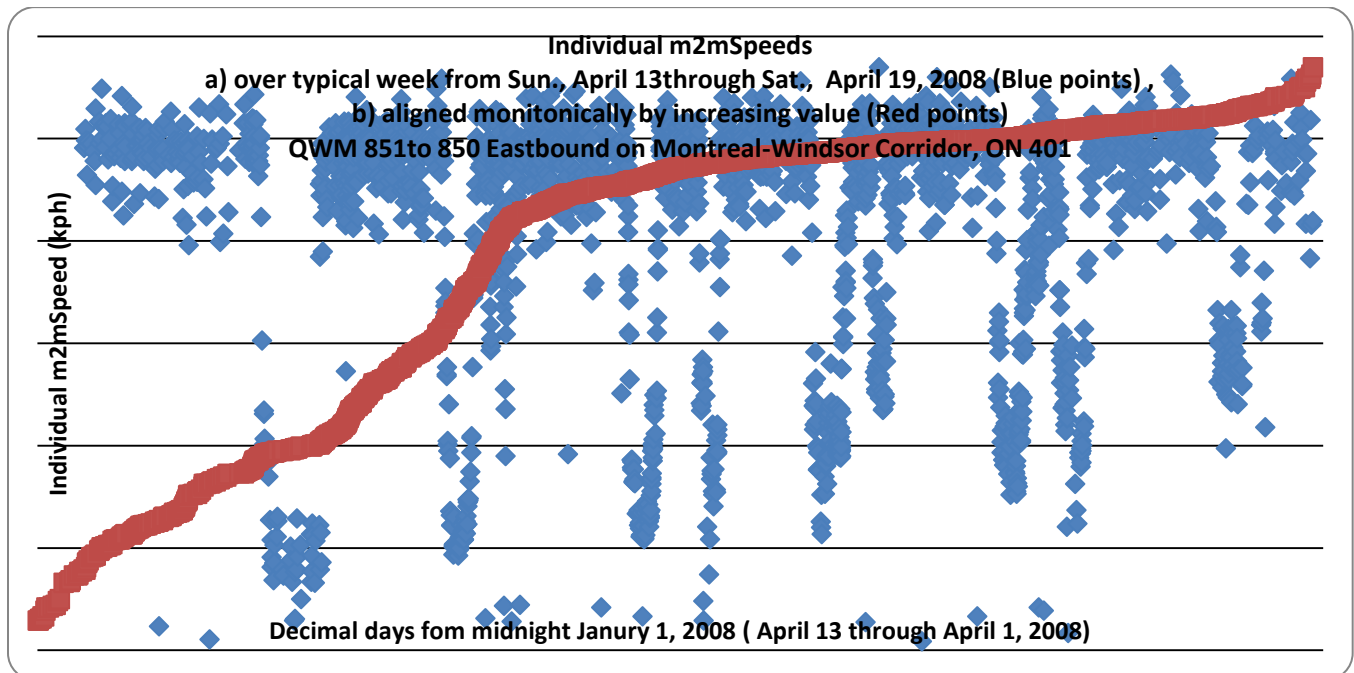

Figure 29, Individual m2mSpeeds distributed by time-of-day and by increasing value for the week beginning April 13, 2008 eastbound on QWM 851 to 850.

The speed performance can be summarized for each corridor, for each direction by computing the median m2mSpeed for any segment for any hour during any of the 8 days of the week, Sunday through Saturday plus Holidays. A three dimensional

surface plot of the Median m2mSpeed summarizes the corridor's performance as is shown in Figure 30 for the Windsor to Montréal corridor. As can be readily seen, most of the corridor is congestion free; however, a couple of segments exhibit substantial recurring congestion (the "valleys" in the surface). In sum, these valleys contributed relatively little delay, on average, but contributed annoying delay to those that happen to be traveling in those locations at those times.
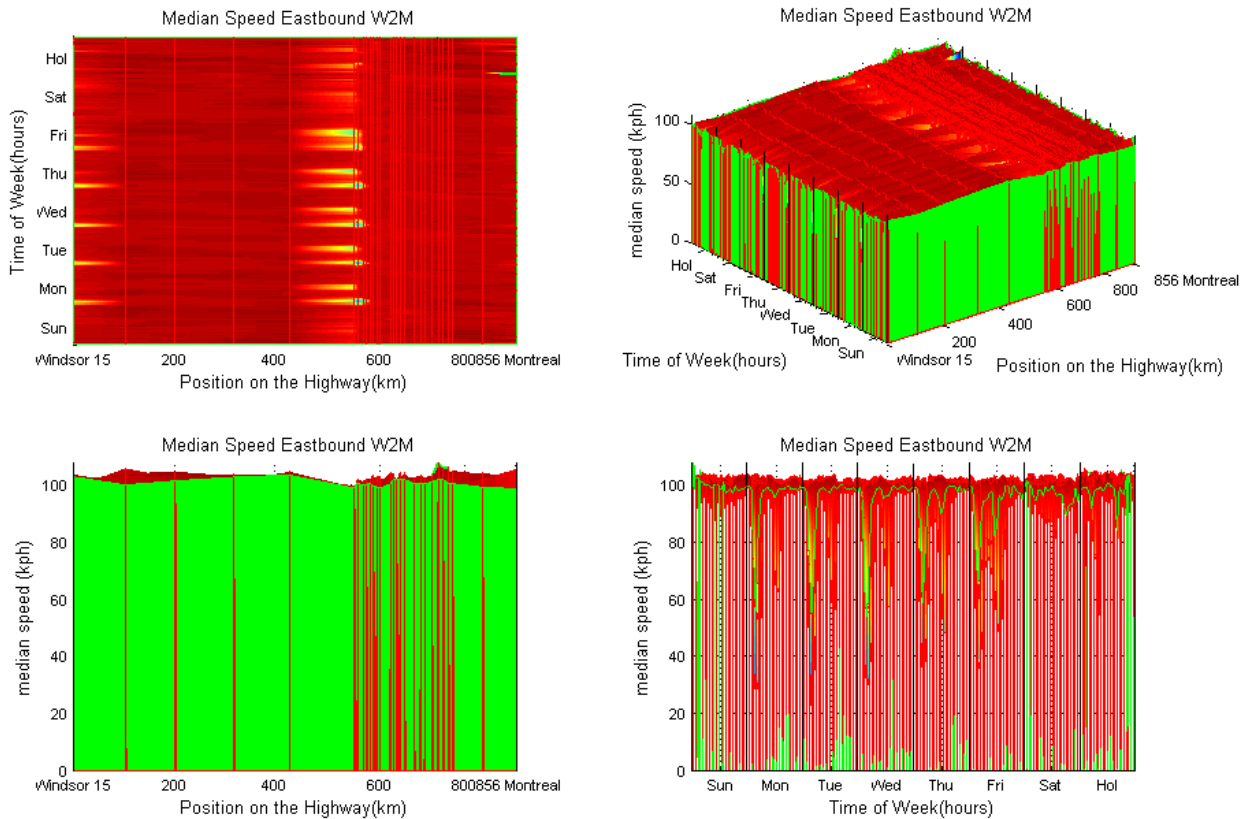

Figure 30, Windsor to Montréal: Median of m2mSpeed each hour by day-of-week across the corridor.

6.0 Example using GPS probe vehicle position data to compute the fastest route from anywhere to anywhere in North America

Since May 1, 2000, when President Clinton signed the executive order that removed "selective availability" that scrambled the publically available GPS signals, ALK Technologies has been harvesting and archiving GPS data captured and shared by users of its CoPilot route guidance system. The data has helped the debugging of the CoPilot software as well as align and extend ALK's North American digital map database. The data has also been useful in investigating route choice characteristics of truckers, Kornhauser, Knorring, and He (2004). With the proliferation of the CoPilot route guidance system, the volume and spatial diversity of the archived data has grown to such an extent that it has become practical to use these data to characterize the speed level-of-service for every road represented in the ALK digital map database of North America, all 38 million arcs representing 7 million miles of road. While some are so lightly traveled that no CoPilot observations are or may ever be available, others have sufficiently high volumes that time-of-day and day-of-week variations can be characterized. The valuable element of the CoPilot sourced GPS data set is its logged data rate. Most tracks have NMEA RMC messages logged every 3 seconds. This high data frequency allows for the characterization of the speed level-of-service on segments that are as short as 100 meters.   The 3 second data rate is many orders of magnitude better than the "every 45 minutes" typical data rate for the QASPR system. Moreover, each RMC message also contains the instantaneous speed and heading of the probe, SiRF (2005). There is no reason to suspect that CoPilot users aren't representative of the travelling public. Thus, it is assumed that the speed distributions of the CoPilot GPS data are a good reflection of the nominal speed level-of-service achieved on any road segment as long as some minimal number of arguably independent observation age contained in the GPS probe vehicle data set.

27

For the current compilation, at least 5 independent speed observations are required to assign a unique median speed and standard deviation to an arc in a direction. If less than 5 independent observations exist, then the assigned speed and variance is that for travel in the opposite direction (if it has more than 5 independent observations and the segment is not one-way) else, it is assigned a regional median value for that road class. The regional median value is the median of the median values on all arcs of the same road class in that geographic area.

Median values for larger segments made up of a chain of arcs are converted simple sums of the median travel time and their standard deviation.

$$\text{Median Measured Average Speed, } \underline{S} \; = \; \Sigma \, D_k \, / \, \Sigma \, \{ \underline{S}_k \, / D_k \} \; \text{ for all k in the chain} \qquad (5)$$

$$\text{Standard Deviation, } \mathbf{SD} = \; \Sigma \, D_k \, / \, \Sigma \, \{ \, SD_k \, / D_k \, \} \; \text{ for all k in the chain} \qquad (6)$$

Where $\underline{S}_k$ is the median speed on link k (in one direction is k is a two direction link), SDk is the standard deviation of speed on that link in that direction, and D_k is the length of link k.

Figures 31 through 40 depict the median speed, $\underline{S}$, by direction for various views of the ALK digital map database. In each figure median speed is displayed along each segment using right-handed bandwidths (the height of the band to the right of the segment is proportional to the median value heading in the facing direction). Colors also indicate speed. The greens are for speeds of 40 mph and above, with bright green being the highest. Reds indicate speeds under 40mph with the slowest being bright red and those closer to 40 mph are light brown.

Figure 31 shows only the major roadways throughout North America. Note that outside the major metropolitan areas, most segments are bright green, meaning that the median speed is substantially greater than 40 mph. Figure 32 is a zoomed- in portion of those major roads near Columbus, OH. Figure 33 displays the speeds on all of the major and minor highways in the same area. Figure 34 is a more zoomed in view of the same class of roadways. Figure 35 shows the highways and streets in the small area near the intersection of I-71 and I-690. Figure 36 shows the speeds only on the major highways while Figure 37 shows the speeds on all road segments. Figure 38 shows median speed on the major and minor roadways near Seattle WA and Figure 39 shows median speeds on all roadways and streets near the center of Seattle. These figures serve to illustrate that median speed as derived from ALK's archived CoPilot GPS tracks have been assigned to each arc in ALK's digital map database of North America. Standard deviation of speed has also been assigned to each link in each direction such that a risk measure can be assigned to a chain of links or any any route, using equation (6).

Since median speed is assigned to every link in the network, the fastest route between any two points on the network can be readily computed and the standard deviation of that or any other route can also be computed. Figure 40 highlights the fastest route from 24 Montadale Circle, Princeton, NJ to the corner of Haight and Ashbury in San Francisco as computed based on the assigned median speeds. The fact that it is very similar to best practical route provides some comfort that the speed assignment process may be working properly. Unfortunately, there exists no proof of correctness. Only the testing of a multitude of routes can begin to certify the process.

7.0 Conclusion

It has been shown that the travel speed on any roadway segment can be effectively characterized using GPS-based probe vehicle position data. The GPS data is of sufficiently high precision and its increasing availability makes it a practical source for monitoring speed on all roadway segments. One way to effectively perform the monitoring is to correlate (map-match) the archived GPS data to a digital map database as a means by which the point-wise GPS data can be transformed into segment-wise information. This process and its result was described and proved effective in monitoring the performance throughout the Quebec –Windsor highway corridor. It is also proving to be effective in monitoring the speed performance of all roadway segments for the purpose of computing minimum expected-time-of-arrival (minETA) routes throughout North America.

Figure 31, Display of Median Speed along the major highways in North America, as observed from ALK's GPS probe vehicle dataset.
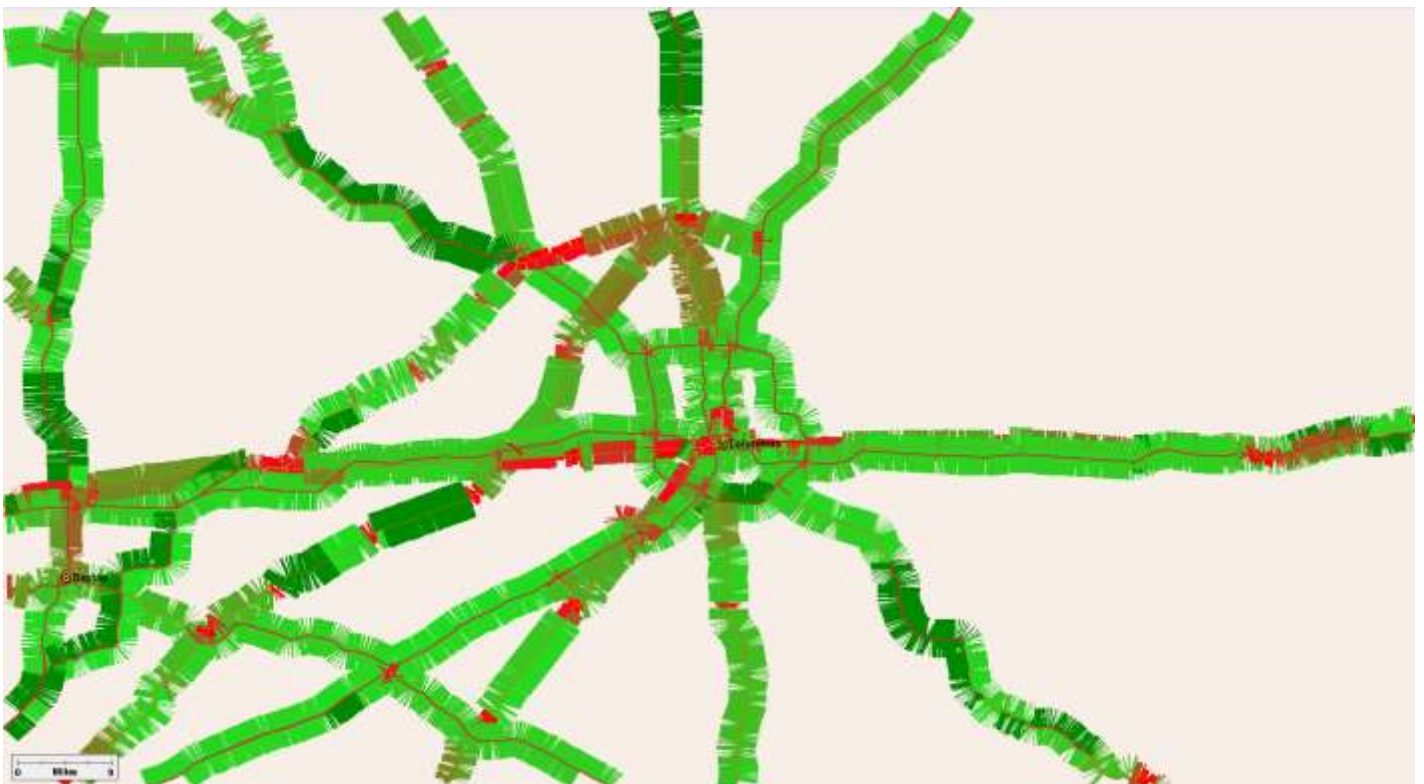

Figure 32, Display of Median Speed along the major highways near Columbus, OH, as observed from ALK's GPS probe vehicle dataset.
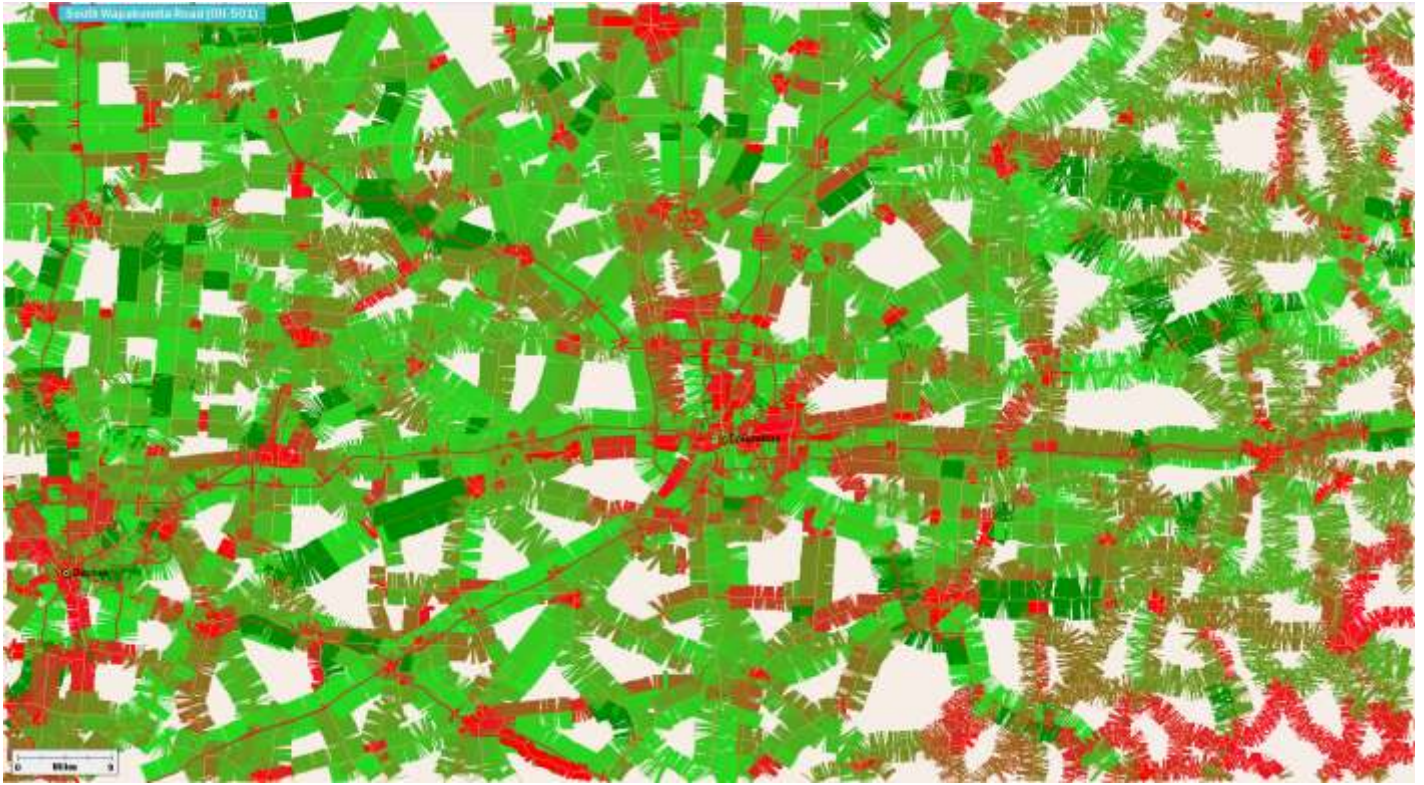
29

Figure 33, Display of Median Speed along the major and minor highways near Columbus, OH, as observed from ALK's GPS probe vehicle dataset.

Figure 34, A more zoomed in display of Median Speed along the major and minor roadways near Columbus, OH, as observed from ALK's GPS probe vehicle dataset.
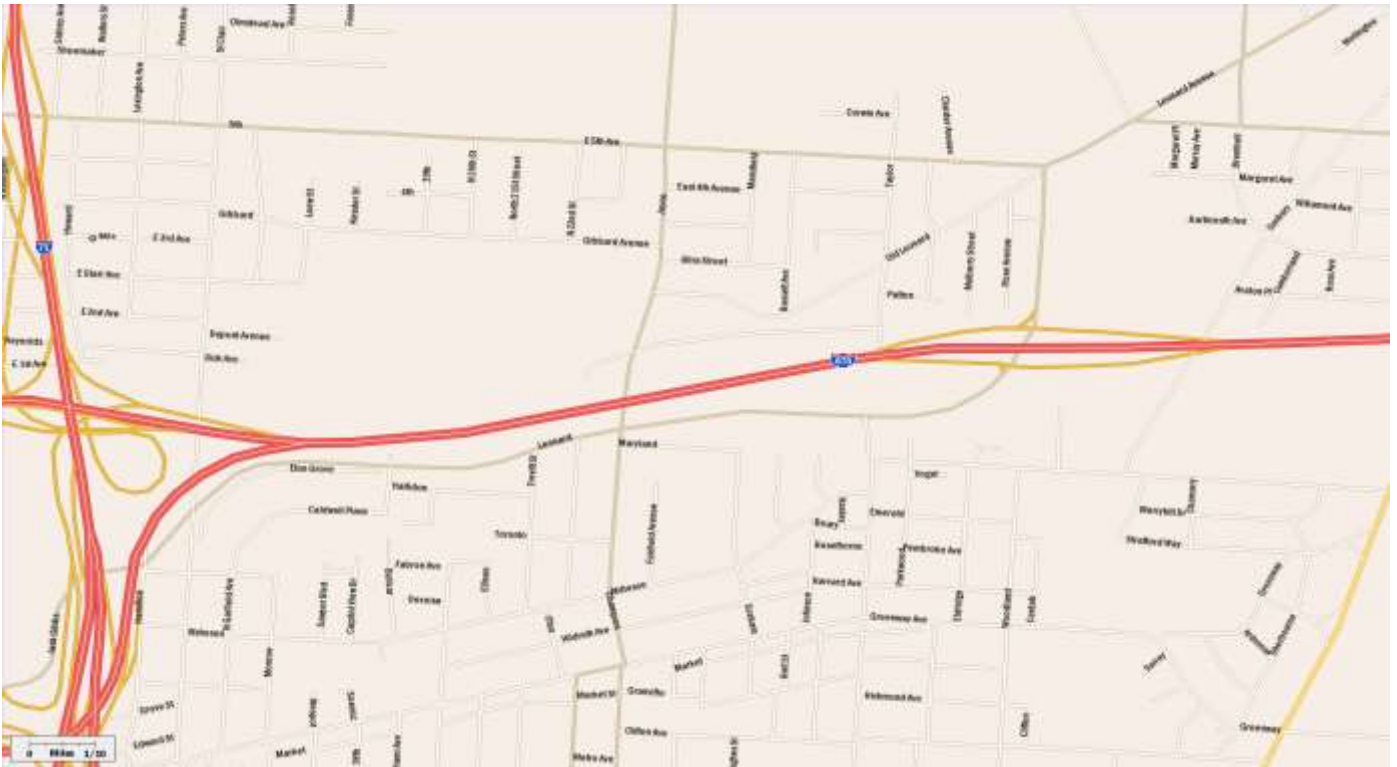
Figure 35, Display of the highways and streets near the intersection of I-71 and I-670 Northeast of Columbus, OH.

Figure 36, Display of Median Speed along the major highways near the intersection of I-71 and I-670 Northeast of Columbus, OH as observed from ALK's GPS probe vehicle dataset.

31

Figure 37, Display of Median Speed along all streets near the intersection of I-71 and I-670 Northeast of Columbus, OH as observed from ALK's GPS probe vehicle dataset

Figure 38, Display of Median Speed along all major and minor roadways near Seattle, WA as observed from ALK's GPS probe vehicle dataset
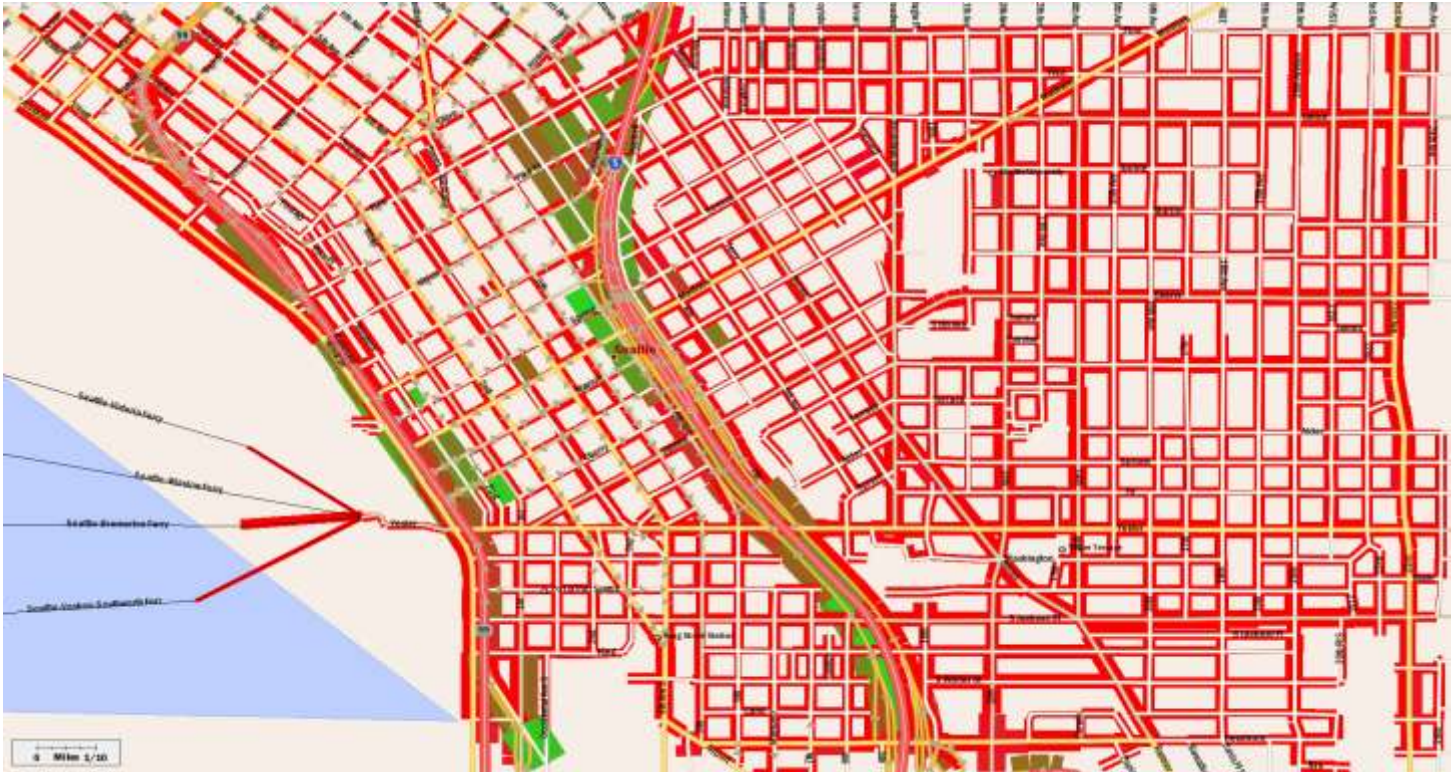
Figure 39, Display of Median Speed along all streets near the center of Seattle, WA as observed from ALK's GPS probe vehicle dataset.

Figure 40, Display of fastest route from 24 Montadale Circle Princeton,, NJ to the corner of Haight and Ashbury in San Francisco based on ALK's GPS probe vehicle based Median Speeds on all links.

8.0 References

Bernstein D., Kornhauser A., (1996) "*An introduction to map matching for personal navigation assistants*" http://www.njtude.org/reports/mapmatchintro.pdf.

Bernstein, D., Kornhauser, A., (1998) "*Map matching for personal navigation assistants"* In proceedings of the 77th annual meeting of the Transportation Research Board, 11-15 January, Washington D.C.

Kornhauser, A. L., U. Batchu (2009). "Statistical Inference Analysis of GPS data to Monitor Highway Performance" Final Report Contract No T8080-08-0122 Transport Canada, Feb., 2009

Kornhauser, A. L., J. Knorring and R. He (2005). "An Analysis of Route Choice by Long-Haul Truckers." Transportation Research Record **1923**: pp. 46-60.

Kornhauser, A. L, I. Lin, and R. He (2008) "Estimating Nationwide link Speed Distributions Using Probe Position Data" Journal of Intelligent Transportation Systems, 12(1):29-37, 2008.

Kornhauser, A.L., and M. Merle (2000)"Truck in the Big Apple: Preliminary Analysis of motor Freight Movements in Manhattan" Presented at Transportation Research Board Annual Meeting, 2001

Kornhauser, A. L., C. Schrader and L. Friese (2004) "Using Historical Information in Forecasting Travel Times" Transportation Research Board annual meeting 2004 Preprint CD

Kornhauser, A. L, et. al. (2006) "Experimenting with Real-Time ATIS: Stepping Forward from ADVANCE" Proc. 9[th] Int. Conf. on Applications of Advanced Technology in Transportation, pp 325-330, ASCE

Morris, A, and A.L. Kornhauser (1999) "Getting Goods Delivered in Dense Urban Areas: A Snapshot of the Last Link of the Supply Chain" Transportation Research Record **1653**: pp. 34-41.

Quddus, M.A. (2006) "High Integrity Map Matching Algorithms for Advanced Transport Telematics Applications" Ph.D. thesis, Centre for Transport Studies Department of Civil and Environmental Engineering, Imperial College London, UK, http://www.cts.cv.imperial.ac.uk/documents/theses/QuddusPhD.pdf

White, C.E., Bernstein, D., Kornhauser, A.L. (2000) *Some map matching algorithms for personal navigation assistants*. Transportation Research Part C 8, 91-108.

Yang, J., K. S. Kang, K. Chon, (2005) "The Map Matching Algorithm of GPS data with relatively long Polling Time Intervals" J. Eastern Asia Society for Transportation Studies, Vol. 6, pp. 2561 – 2573 *personal navigation assistants*. Transportation Research Part C 8, 91-108.

SiRF (2005) NMEA Reference Manual, Jan 2005 http://www.sparkfun.com/datasheets/GPS/NMEA%20Reference%20Manual1.pdf