READING BLUETOOTH JOURNEY TIME MONITORING SYSTEM

PRACTICAL APPLICATION AND LESSONS LEARNT

Simon Beasley: Reading Borough Council, Rob McDonald: Peter Brett Associates LLP

Introduction

A significant number of vehicles carry Bluetooth devices in the form of sat-navs or hands free phone systems. This then, provides the opportunity to collect journey times for vehicles by placing Bluetooth sensors at strategic points throughout a road network. Reading Borough Council have recently completed installation of a TDC Bluetooth Journey Time Monitoring System, with data storage, matching and filtering provided by Drakewell. There are over 130 detectors placed on a variety of road types and these are able to provide real time information on journey times across the network. This paper discusses the benefits of using such a system as well as the lessons learnt and describes the validation work undertaken by Peter Brett Associates (PBA) to ensure that the system is reporting accurate data.

Figure 1: Bluetooth detector on a signal pole

The Bluetooth system

Reading Borough Council purchased a Bluetooth journey time monitoring system with the purpose of improving UTMC network management strategies as well as providing traveller information through open data (http://opendata.reading-travelinfo.co.uk/), variable message signs, website etc. This system is capable of responding in real time as the Bluetooth sites report their data based on a rolling 15 minute average updated every 5 minutes, and therefore it is possible to identify increasing congestion and as such advise motorists on alternative routes.

The systems works through a series of roadside detectors which record the unique MAC address of any transmitting device and then sends this back to the Drakewell C2 system. By matching MAC addresses the system can then output link journey times. The detectors have been predominately installed on traffic signals, using spare signal cores to provide power. Where sites were chosen without traffic signals the devices were installed on solar powered poles and were often combined with new ATC sites. Lamp columns were not an acceptable solution for RBC, but have been used where detectors have been installed by the business parks to supplement RBC's system and provide information for their occupants.

Communications were provided by 3G sim card for all sites as this was more cost effective than cabling back to the broadband services in the traffic signal controller. This was fine for the majority of the sites. However some sites with poorer connections have been prone to delays in the data being received and although all the data is in the system at the end of the day; this can impact on the real time data quality. A number of sites were found to have faulty sims leading to delays in data being sent. The sims were replaced at these sites to ensure that data is uploaded in real time.

At the time of validation there were 135 Bluetooth locations in the Reading urban area of which 96 were used to perform the validation. The reason for this lower figure at the time of validation was due to not all the sites being fully operational and a key factor here was the LSTF programme of traffic signal upgrades which has only recently completed, but there were also other factors including communications issues.

Figure 2 provides a map of the Bluetooth locations as well as showing which sites were used for validation purposes.

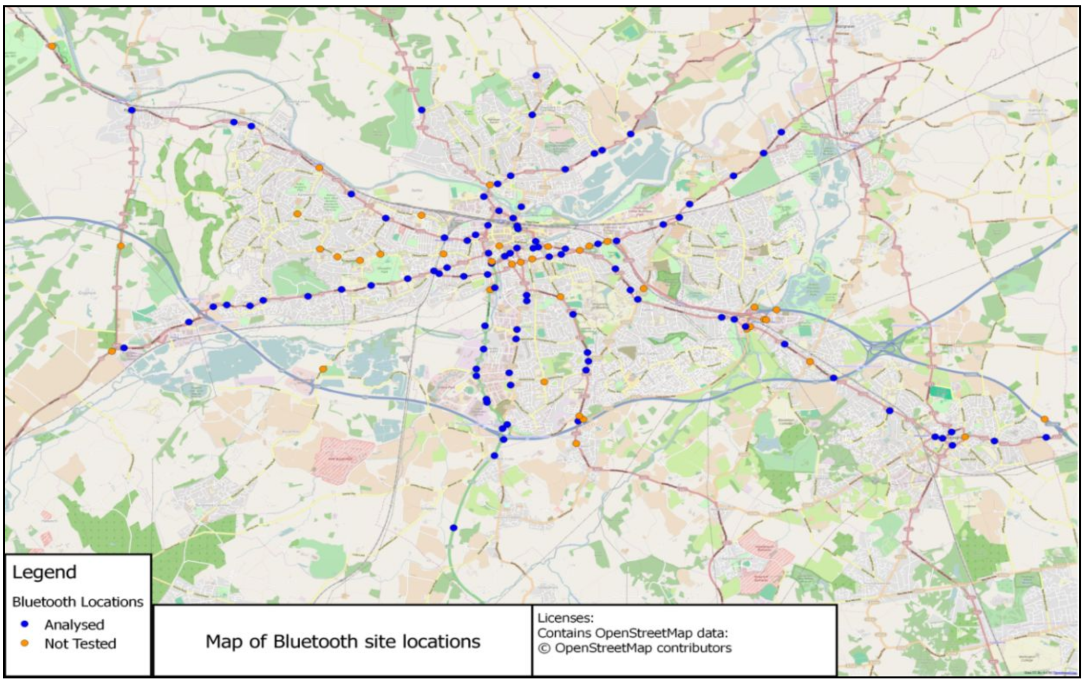

Figure 2: Map of Bluetooth sites

The sample rates at any individual site were found to be as expected from other installations with between 20% and 45% of the vehicles passing any one device having Bluetooth enabled.

The output data is fed into a hosted Drakewell C2 system that provides data filtering (automatically removing the outliers on a rolling 15 minute basis) as well as providing the interface that allows end users to view the data and generate reports. This 15 minute rolling approach of updating journey times ensures reliability in the information if counts are low in any 5 minute period, however it does also dampen the ability of the system to respond to a real incident. It is understood that Drakewell are working on this to improve reporting and allow shorter reporting periods where counts are high and a more weighted approach to more recent data which would improve the real time responsiveness of the system.

Methods of calculating journey times

The Drakewell system offers two different ways to analyse the data for any given route. The first system is to create routes joining a number of Bluetooth sites together. The reported journey times from this feature are generated by summing the journey times for each link along the route. This system is also able to compensate for sites that are offline by skipping over a site and matching MAC addresses for the next shortest link that is available.

The second method is to simply pick the sites at either end of a route and evaluate journey times between the two however this then only uses matches for vehicles making the whole end to end route. The routes generated for the validation are shown in Figure 3.

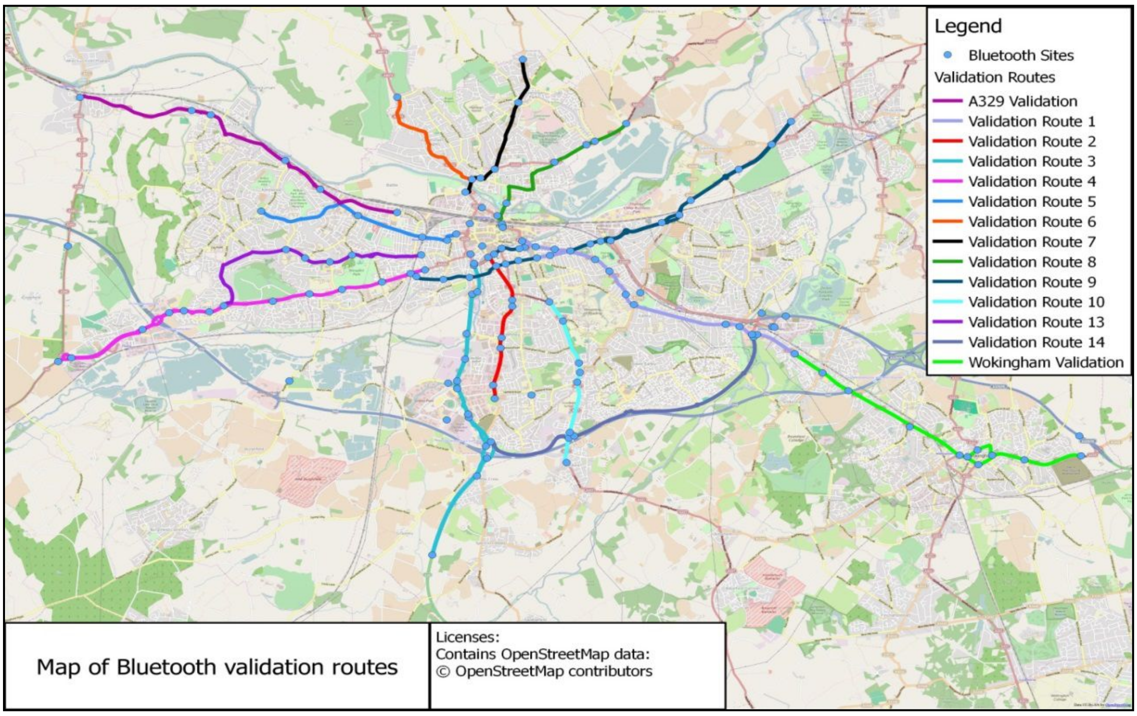

Figure 3: Map of validation routes

Possible limitations of using Bluetooth data as a source of journey times

The primary data issue is that Bluetooth detectors will pick up Bluetooth signals from sources other than cars (such as mobile phones belonging to pedestrians or cyclists etc.). This has the potential to lead to the journey time data being skewed and not being reflective of the current traffic conditions. The Drakewell system automatically does some filtering and the main scope of this article is to discuss how accurate the data is following the filtering process. The filtering consists of removing any data that is more than 4 times the 25th percentile, then another layer of filtering uses the spread of the data to remove any potential outliers based on how far each data point is from the rest of the data.

Another possible issue with the data was clustering of data. For example a bus full of passengers may have a number of active Bluetooth devices which could bias the journey time towards the bus journey time rather than that of the traffic. On links with bus stops this could lead to slower journey times and on links with bus lanes it could lead to faster journey times. It should be noted that over 90% of buses in Reading are fitted with AVL and hence the Bluetooth system was not implemented to try to generate bus journey times.

Validation of the Reading system

The validation process has been undertaken by comparing the Bluetooth journey times between the sites to GPS (Floating Car) data provided by Trafficmaster. This gave an insight into whether or not the Bluetooth system accurately predicts journey times against the observed GPS data. The Floating Car data contains more Light Commercial Vehicles (LGVs) than normal traffic (57.8% of the data set) however this isn't likely to have a significant impact on the comparison as journey times of an LGV are likely to be similar to that of a car in the urban area where congestion is the key determinant of journey time.

It was decided to perform the validation on the morning and evening weekday peaks as these are the most important times for traffic management. The data has been collected in 15 minute intervals for each link between sites as well as an end to end summary for each route as a whole. Note that 15 minute intervals were chosen over 5 minute intervals (the frequency that the COMET Urban Traffic Management and Control System updates) as the validation data was supplied in 15 minute intervals.

As well as checking the data accuracy on each link in the network a comparison was made to look at the correlation between the Bluetooth and GPS data sets. This can also provide another comparison between using link by link data against end to end route data. A linear regression model was also created to see how well the Bluetooth journey times predicted the GPS times. This in particular provided good statistical evidence of how accurate the Bluetooth data is compared to the GPS data. The regression identifies whether there is a bias in the system (i.e. if the system is routinely out by a certain number of seconds) as well as showing whether the Bluetooth system under or over predicts journey time. If the system is working perfectly the expected regression fit would have a slope of one (i.e. direct relationship between Bluetooth and GPS) and an intercept of zero.

The analysis of data clustering comprised of looking at the data set for each link and looking for any data points where the start and end points are within 5 seconds of another recorded journey. These points are then removed and then the results are compared against the raw data. A 5 second period is quite large and may filter out some trips where vehicles are close together but it will allow for filtering out clustering in longer vehicles such as a bus. This filtering and analysis was also done for 1, 2, 3 and 4 seconds as the cluster threshold but the results for five seconds are reported as they represent a worst case. Additionally links that contain bus lanes were analysed in more detail to see if there was any evidence of bus passengers affecting the reported journey time as any such passengers will likely appear in a cluster in the Bluetooth system.

The validation process also included analysis into whether the mean or median should be used when reporting journey times. The median has the advantage that any points that are far from the norm have less weight and therefore less likely to skew the data. The mean however has the advantage that all of the data is used and the data in theory should already have been filtered for outliers.

Results

In total there were 160 links analysed with 80 inbound links and 80 outbound links. A link was determined to be accurate if it was within 20% of the GPS time. This is to allow for random fluctuations in the GPS and Bluetooth time. This extends to other factors such as variability caused in wait time at traffic signals between detection points as well as variation in detection at sites which will depend on when each Bluetooth device was detected. The total number of links that were not within the threshold described was 15. In many cases this was due to apparent outliers in the floating car data used as a comparison (which is due to low sample rates on less frequent routes for any give 15 minute period). The analysis was redone on such routes using a month's worth of data instead of two weeks and in most cases the Bluetooth links were found to be within the 20% threshold. Links that were reporting journey times beyond the 20% threshold tend to be shorter routes which are much more vulnerable to variations in detection time and delay caused by traffic signals as well as the fact that a 20% difference on a link of 30 seconds is very small in absolute terms (6 seconds).

The results of comparing the correlations between the data sets suggest that the Bluetooth data is a good representation of the observed floating car data. The correlation was 92.70% (using the mean journey times) or 92.63% if the median is used. These values come from looking at each link individually. When matching sites at either end of a route the correlation falls to only 33.93% (or 33.53% if the median is used). This then suggests that links should be kept relatively short with this kind of data as the distance between end points of a route has the potential to cause accuracy issues. The reason for these accuracy issues is that end to end matching is low and too high a proportion of vehicles stop, take alternative routes etc., so that there is not enough data for the automatic outlier removal to recognise outliers.

Alternatively when route data is reported as the sum of each link it comprises, the correlation is 97.27% between Bluetooth and GPS. Which means this method should be preferred over using the end points only for a route.

The results suggest a trade-off is needed between link length and amount of data available. The end to end links are inaccurate due to small sample sizes as many vehicles turn off before the end of the route. Short links are likely to be affected by variations in signal timings and detection by the sites. The shortest links are around 200m and as low as 100m on the Inner Distribution Road (IDR). Links appear to be at their most accurate when the link length is over 300m and where there is sufficient data to perform filtering.

The regression analysis for the individual link data results in an intercept coefficient of 6.22 and a slope coefficient of 0.93. This would result in a line of best fit of where x refers to Bluetooth time and y refers to GPS time. The value of the intercept is fairly low at less than ten seconds and the slope is fairly close to one. This demonstrates that the Bluetooth system is performing as intended. Note that these results are obtained using the median journey time rather than the mean.

Figure 4 shows the regression plot comparing the Bluetooth and GPS data. The "Theoretical" line is the relationship you would expect if the Bluetooth data was a perfect match for the GPS data. The figure clearly shows that the Bluetooth data is very similar to the GPS as the regression line is very close to the theoretical line.

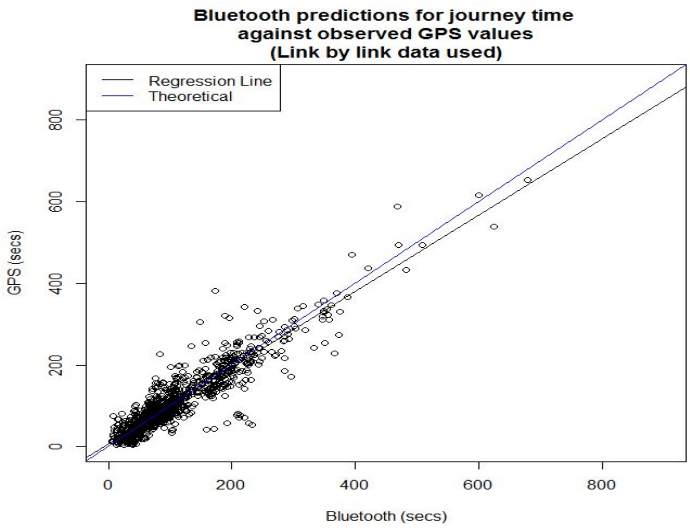

Figure 4: Regression plot of Bluetooth against GPS

The same process was not undertaken for the routes as a whole as the data sets are too different to attempt to plot a linear relationship between the two.

Clustering results

The data showed that most of the links are affected by less than 1% by clustering (70.28% of the link by link data and 63.54% of the end to end route data). Only 0.71% of the link by link time periods have a clustering effect of more than 5% (8.33% for the end to end data).

This shows that matching across whole routes is more prone to being affected by data clustering (likely due to the lower number of matches due to the increased distance). The data shows that individual links are not significantly affected by clustering. This was further investigated by looking in detail at the links with higher clustering to see if there was a pattern – i.e. it was those with bus lanes but this was not found to be the case with no clear correlation.

Mean against Median

The results showed that there isn't a significant difference between using the mean or median. This can be seen by the similar correlation between the Bluetooth and floating car data sets when the mean and median are used (92.70% to 92.63% respectively). The results of the regression analysis are also similar with a fit of when median data is used which is very close to from the mean data. It should be noted that the median is the preferred method as it offers greater resilience against outliers than the mean does.

An interesting anomaly

In the initial installation a 'super site' was inadvertently created on the A33 dual carriageway where a Bluetooth device had been installed on each carriageway but the detection zones were overlapping. Hence the same vehicle was being detected at the same time by both northbound and southbound detectors and was then being matched. The journey distance for this movement involved U-turning around a roundabout and hence the system generated very high average speeds in the order of 150mph.

Summary

The analysis shows that Reading's Bluetooth system offers a reliable way of measuring journey time. It has also shown that when planning a new system care should be taken in locating detectors to optimise the results achieved. The findings of the analysis indicate that links should, wherever possible, be in excess of 300m while balancing this against the need to minimise the chances of vehicles exiting the link and therefore reducing the capture rates. There is therefore also a balance to be made in the ability to provide origin / destination data, which requires detectors to be located at key intersections in the network, with the reduced accuracy of journey time data when this reduces link lengths below 300m.