Oyster-Based Performance Metrics for the London Overground

Michael Frumin

Strategy Team, London Rail, Transport for London
School of Civil and Environmental Engineering, MIT
mfrumin@mit.edu

October 28, 2008

Executive Summary

The London Overground is a pre-existing rail service in London whose operating responsibility and revenue risk were recently granted to Transport for London (TfL). Here we discuss the prospect of using data from the Oyster smartcard ticketing system to evaluate the performance of the London Overground explicitly from a passenger's perspective.

The core idea behind our approach is to directly measure end-to-end individual journey times by taking the difference between entry and exit transactions stored by the Oyster system. The focus of this study is *Excess Journey Time* (EJT), calculated on a trip-by-trip basis as the difference between the observed journey time and some standard. In this case, the standard is determined for each trip with reference to published timetables, indicating how long the trip should have taken under right-time operations. A positive EJT indicates that the journey took longer than was expected.

Excess Journey Time is interpreted as the delay experienced by passengers *as a result* of services not running precisely to schedule. The distribution of EJT indicates reliability. We validate these interpretations using a detailed graphical analysis, and then aggregate them to the line and network level over a variety of time periods. Our analysis is conducted on large samples of Oyster data covering several months and millions of Overground trips in 2007.

At the aggregate level, relative values of Excess Journey Time are largely in line with expectations. The North London Line has the highest average Excess Journey Time of all lines on the London Overground, around 3 minutes, and the widest distributions (i.e. least passenger reliability). On all lines, there is significant day-to-day variability of Excess Journey Time. For the whole London Overground, and for the North London Line in particular, Excess Journey Time is worst in the AM and PM Peak timebands.

The current performance regime for the London Overground is the Public Performance Measure (PPM), which measures the fraction of scheduled vehicle trips arriving at their destinations fewer than five minutes late. Over time, EJT shows a strong correlation to PPM. There is clear additional variation in EJT, indicating that it captures certain information about passenger experiences that PPM does not. This variation tends to increase as PPM decreases, particularly in the AM and PM peak timebands, which suggests that the effectiveness of PPM as a measure of the passenger experience decreases as service deteriorates.

Another quantity of interest derivable from Oyster data is the time between passenger arrival at the station and the scheduled departure of the following train. The spread of this distribution of this quantity indicates the degree to which passengers arrive randomly (i.e. "turn up and go") rather than time their arrivals according to schedules. We have found that on the North London Line, especially during the AM, interpeak, and PM peak periods, passengers tend to arrive randomly. This is apparently in contrast to conventional wisdom for National Rail services, and has distinct implications for crowding levels and timetabling practice. In an appendix to this report we look at this in detail, and recommend that even headways be prioritized in timetabling the North London Line.

The Overground is, by design, part of a larger integrated multimodal network. Oyster data, by nature, is somewhat ambiguous in representing passenger trips on such a network that

involve transfers or multiple routing options. This poses certain problems to our methodology, but also presents the opportunity to quantify and understand the experience of passengers across the entire network. We discuss these problems, potential solutions, and opportunities at length, as well as other applications for this methodology, and future research directions.

We have concluded that Oyster-based metrics are effective for monitoring and identifying problems as experienced by passengers on the London Overground. They may be even more effective for use across the whole of London's public transport network, particularly as Oyster is in the process of being rolled out to all National Rail services in the Greater London Area.

Introduction

The London Overground is a pre-existing rail service in London whose operating responsibility and revenue risk were recently granted to Transport for London (TfL). Here we discuss the prospect of using data from the Oyster smartcard ticketing system to evaluate the performance of the London Overground explicitly from a passenger's perspective.

The core idea behind our approach is to directly measure end-to-end individual journey times from the entry and exit transactions necessitated by the zonal fare structure of London's transport system. Comparison of these observed journey times to expected journey times quantifies the delays experienced by passengers. Because the Oyster system provides extremely large sample sizes we can also measure reliability as indicated by distributions of journey times, which we believe is fundamental to understanding how the railway's operational performance affects passengers.

By virtue of the Overground's National Rail origins, the regime used to evaluate its operations remains in line with that used by other National Rail services. That system, known as the Public Performance Measure (PPM), describes operations in terms of the punctuality of scheduled vehicle trips. This is out of line with other metrics used by Transport for London's Bus and Underground services, which are expressed in units of weighted passenger time.

Our intent is not to replace the PPM, but to complement it with a deeper and more nuanced appreciation of the passenger experience than is currently provided. The PPM describes the *punctuality* and *performance* of scheduled vehicles; we add measures of *delay* and *reliability* for passengers.

The work detailed in this report was conducted in June and July of 2008 in TfL London Rail as part of the ongoing research collaboration between the Massachusetts Institute of Technology (MIT) and Transport for London. This proof-of-concept project developed a methodology, anticipated problems, developed strategies for wider application, and produced preliminary results.

This report is structured as follows. Part 1 provides some context by describing a number of performance regimes already in use at TfL. Parts 2 and 3 describe the data and methodology used to measure the performance of the Overground. Part 4 presents an exploratory analysis of Overground travel patterns, and in Part 5 the approach is validated via graphical analysis. Part 6 presents aggregate measures of passenger delay and reliability, and Part 7 relates these to the metrics produced by the current performance regime through linear regressions. Part 8 discusses a number of likely problems, and potential solutions. Parts 9 and 10 present the prospects for future use and further research. Finally, we include an appendix examining the effects of turn-up-and-go passenger arrivals on vehicle crowding as an example of other types of analysis that are possible with Oyster data.

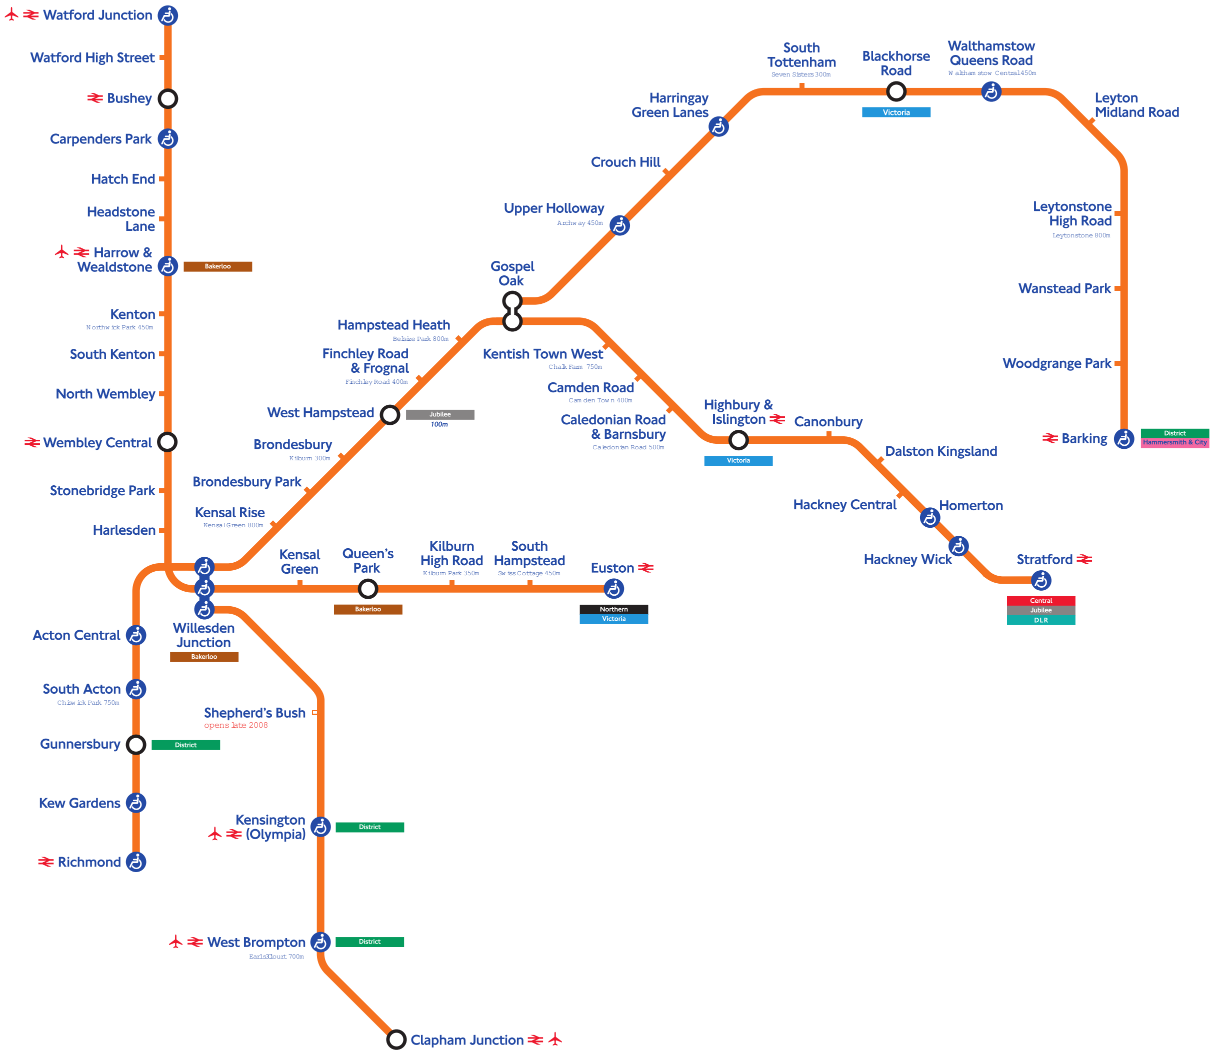

Figure 1: London Overground Network Map

1 Existing TfL Performance Regimes

Transport for London uses a range of metrics and indicators to measure the performance of the services it provides. Three of those -- the London Overground's Public Performance Measure (PPM), London Underground's Journey Time Metric (JTM), and London Buses' Excess Waiting Time (EWT) -- are relevant to this study and are briefly summarised below.

Public Performance Measure – London Overground

As mentioned above, the PPM measures performance of the London Overground, as well as all other National Rail passenger services. The key PPM metric is the fraction of trains that have arrived at their destination less than 5 minutes behind schedule, which is sometimes disaggregated to the line and/or timeband level. Another metric, the so-called "Delay Minutes," measures the total delay incurred by scheduled vehicles along their route. These measures and others are generated, with culpability assigned to a variety of sources, using Network Rail's "TRUST" database system, which logs (with much human effort) structured and unstructured information about every scheduled vehicle trip and every incident on the network (ORR 2008b).

A recent sample of the London Overground's PPM report is included in Appendix A. More information on PPM can be found on the *Performance and Financial Results* page of the Network Rail web site (NR 2008a).

Journey Time Metric (JTM) – London Underground

One of the London Underground's primary measurements of performance from a passenger's perspective is the Journey Time Metric. The JTM measures the times of each component of a trip – Access/Egress/Interchange, Time in Queue (for tickets), Platform Waiting Time, and On Train Time – using a variety of data sources, some manually sampled and some automatically generated, but none taken from direct end-to-end observation of passenger journeys. For example, egress times from platform to station exit are measured by surveyors with stop watches, while on train time and platform waiting time are modelled with data from train signalling systems. These kinds of measurements are taken numerous times over the entire London Underground network over each four week period.

In all cases, measured journey times for each trip component are compared to scheduled, standard, or free-flow journey times to calculate *excess journey times* (EJT), which are then multiplied by static measures of demand over the network. Excess Journey Time is reported for different lines, facilities, and time periods in weighted and unweighted format. Weighted journey times represent additional time that is perceived by passengers for crowding and for trip components other than in-vehicle travel time.

The Journey Time Metric strives to measure something quiet similar to what we hope to measure with Oyster data, but there are some key differences. The JTM, by nature of its component-wise measurements, directly exposes a more granular view of precisely where in the network passengers experience delay. However, its sum-of-parts approach may not fully capture the often-correlated effects of multiple delays affecting trips through the network, which should be explicitly included in Oyster-based end-to-end measurements. As well, it uses static measurements of demand to estimate how many passengers experience these delays whereas Oyster directly measures demand on a day-to-day basis. Finally, the variance of distributions of journey times (i.e. reliability), rather than just mean values, can be measured from large samples of by Oyster data.

The ongoing collaboration between London Underground and MIT has produced great strides towards using Oyster data to estimate origin-destination matrices and to measure journey times and reliability (Chan, 2007). More information on the Journey Time Metric methodology, and reports of recent results, can be found on the London Underground's corporate intranet (LU 2008).

Excess Waiting Time (EWT) – London Buses

London Buses' primary means of assessing performance from the passenger perspective is the Excess Waiting Time (EWT) metric. Unlike the JTM, the EWT focuses only on the waiting time component of bus journeys. Like the JTM, it does not make use of direct measurement of actual passenger journeys. Instead, it estimates average waiting times on a given route with a mathematical model that relates average waiting time to observed vehicle headways. The difference between this result and the expected waiting time given perfectly even headways is considered the excess waiting time (note: the same mathematical model is used to estimate excess platform waiting times in the JTM). More information on the EWT methodology and recent results can be found on the TfL corporate intranet (LB 2002), (LB 2008).

A program to install GPS systems on all London buses is currently underway, with the explicit purpose of measuring the performance of the various concessionaires who operate London's many bus routes under contract from TfL. These systems will give broad and deep measures visibility into the performance of operations, but will not add explicit information about actual passenger journeys.

2 Data

This section describes the various data sources used to complete this study.

Oyster Journeys

The primary inputs into this study are two different samples of recent data from the Oyster system. Both samples are limited to trips beginning and ending at London Overground stations that were paid for using Oyster (be it on a Travelcard or Pay as You Go). The record for each trip consists of an entry and exit transaction, a small sample of which is shown in Appendix B.

The two samples are:

1. 100% sample of Oyster trips between any two Overground stations during the 50 business weekdays from 31 March to 10 June 2008, inclusive. This amounts to over 2.6 million trips made on over 417 thousand different Oyster cards.

2. For a random 5% panel of all Oyster cards, all trips between any two Overground stations during the 116 business weekdays from 4 February to 18 July 2008, inclusive. This panel is updated with 5% of all new Oyster cards at the end of each 4-week period. It contains over 325 thousand trips made on nearly 28 thousand different Oyster cards.

The advantage of the former sample is that it is a much larger sample, allowing for more disaggregate analysis. The latter, smaller, sample covers a substantially longer span of time.

Each Oyster transaction record contains a number of components describing the stations used, times, dates, payment type, fare, capping, and an anonymized identifier for the card itself. Of primary interest to this study are the stations, dates and times of entry and exit. These times are truncated to the minute, so some accuracy is lost at the individual trip level.

Many of these trips are eventually deemed not to have been made on the Overground and thus excluded from the sample. Likewise, many trips that are likely to have used the Overground as one leg of a multi-part journey are never considered because they either began or ended at a station without Overground services. These and other issues concerning the inherent ambiguity in Oyster data, and strategies for dealing with it, are discussed at length in Section 8 of this paper.

PPM/TRUST Data

Another important input is the TRUST delay data and daily PPM calculations from the start of Overground service on 11 November 2007 to the current. These come in the form of the PPM summary for each line of the Overground for each timeband for each day, and individual records for every single TRUST incident.

Other Inputs

Other inputs include the London Overground timetable (exported from MOIRA), a simple flat-file based representation of the Overground network (built by hand), and various reference tables that come with Oyster data.

3 Methodology

This section describes the methodology developed to measure delay to London Overground passengers using data from the Oyster smartcard system.

Excess Journey Time and Scheduled Waiting Time

Measuring the end-to-end journey time of an individual Oyster trip is as straightforward as taking the difference between the timestamps of the entry and exit transactions. When a trip involves gated stations, the access and/or egress times between gates and platforms are part of this time, though not so when platform validators are used. In all cases, platform waiting time is included, as is any time spent transferring between services on a multi-leg trip.

It is possible that with advanced inference techniques the different components of a trip – access, waiting, travel, transfer, and egress – could be estimated separately and weighted accordingly. That is beyond the scope of this study so all journey times presented here are unweighted. This is not in line with the Underground's Journey Time Metric or conventional calculations of Generalised Journey Time, which assign separate time weights to each of the different components of a trip.

Given that the Overground is a rail (rather than teleportation) service, passenger journey time is of interest but is not in and of itself a problem. Of prime importance, and the focus of this study, is *excess* journey time -- the amount of time a trip takes less the expected time for that trip.

While there is little to debate about how long an observed trip actually took, defining the expected time for that trip is a matter of policy or standards. Here we define the expected, or scheduled, journey time according to the timetable under the following assumptions:

- The journey begins in the minute when the passenger touches his or her Oyster card at the station of origin
- Passengers take the first train which departs from their station of origin (towards their destination) on or after the minute in which their journey began
- The journey ends in the minute when the train arrives at the station of destination

Thus, the *Scheduled Journey Time* (SJT) for a given trip is the time between the Oyster touch-in and the earliest possible arrival at the destination, assuming all trains run perfectly to published timetables. The *Observed Journey Time* (OJT) is the time between the Oyster touch-in and touch-out. The *Excess Journey Time* (EJT) is the difference between the two:

Excess Journey Time (EJT) =
Observed Journey Time (OJT) - Scheduled Journey Time (SJT)

A positive EJT indicates that the journey took longer than expected; a negative EJT indicates the opposite. Because the Observed Journey Time includes access and egress between gates and platforms but Scheduled Journey Time does not, we expect to see a positive Excess Journey Time in some cases even when trains ran precisely to schedule. This is not felt to be a major problem because the distance between gates and platforms on the London Overground is in general relatively short, because Oyster data is accurate only to the minute, and because EJT is likely to be used as a relative measure, tracked over time.

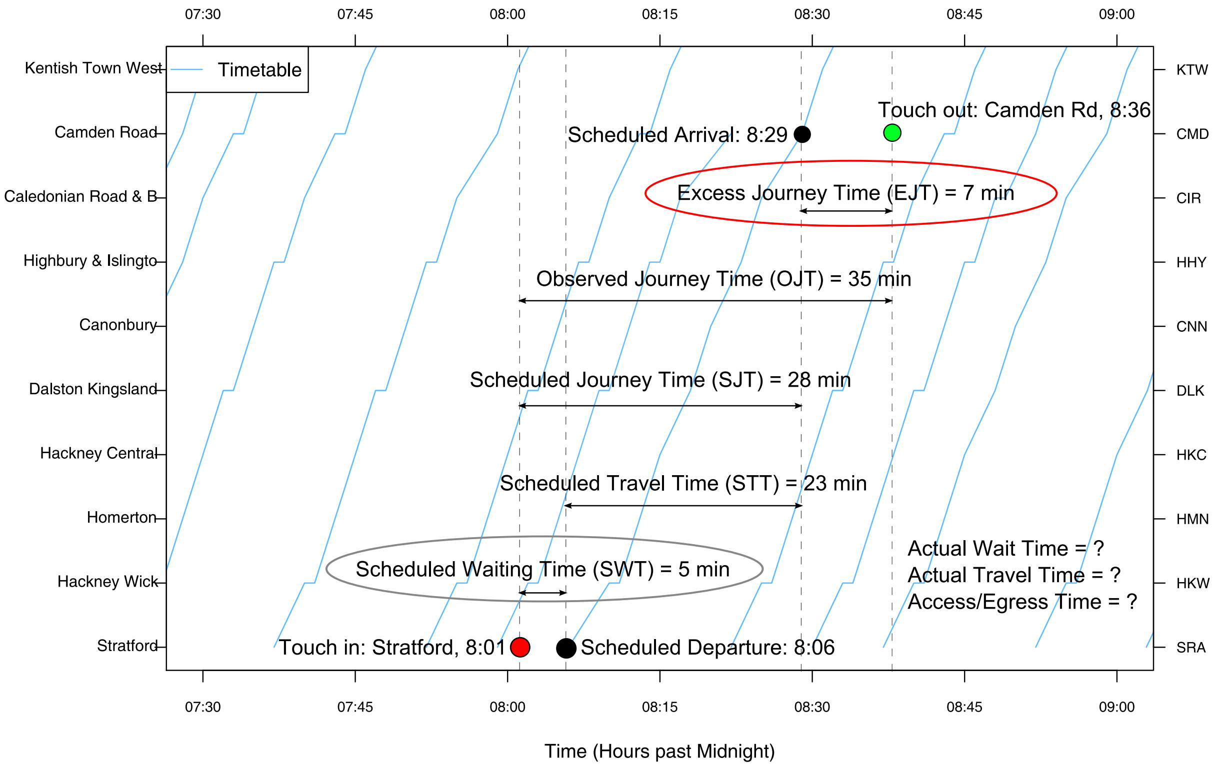

Figure 2: Illustration of calculation of Excess Journey Time

Another quantity of potential interest is the *Scheduled Waiting Time* (SWT) – the time between a passenger's arrival at the station and departure of the following train. Figure 2 illustrates the calculation of EJT and SWT for one example trip from Stratford to Camden Road on the North London Line.

Network Structure and Overground Journeys

Oyster data, by nature, do not come with any indication of which actual services were used to travel between the origin and destination stations. To generate a Scheduled Journey Time for any given trip, we must decide first what route was taken, and then for the chosen route which individual service or services were used.

This process is made quite complex by the extensive rail network in and around London, and is discussed at length in the Ambiguity section. For the sake of simplification, we adopted the scheme used by TfL's Fares & Ticketing department which identifies certain origin-destination pairs as London Overground routes.

By default, trips between two Overground stations are classified as Overground trips, while origin-destination pairs that are interavailable or are better served by the Underground or other National Rail services are classified as such. Examples of these different classifications are shown in Table 1. Again, trips that either start or end at a station without any Overground services are neglected from this analysis entirely.

Origin Destination Type	Example O/D Pairs
London Overground Only (LO)	Stratford ↔ Dalston Kingsland Willesden Junction ↔ London Euston
London Underground Only (LU)	Stratford ↔ Richmond Highbury & Islington ↔ Blackhorse Road
Overground + Underground (LO/LU)	Richmond ↔ Clapham Junction
Interavailable	Richmond ↔ Gunnersbury Wembley Central ↔ Queens Park
London Midland Rail (LMR)	Watford Junction ↔ London Euston Harrow & Wealdstone ↔ London Euston
Overground + London Midland (LO/LMR)	Carpenders Park ↔ London Euston North Wembley ↔ London Euston
Southern	Clapham Junction ↔ West Brompton Clapham Junction ↔ Wembley Central

Table 1: Classifications of Origin Destination Pairs Between Overground Stations

As we will see this classification is not perfect, but it allows us to move forward with relative confidence. The two samples used in this study are of course reduced to the following degree when including only those trips deemed to have used only the Overground:

- 100% sample: Reduced from 2.6 million trips on 417 thousand cards to roughly 1.8 million trips on 308 thousand cards.

- 5% sample: Reduced from 325 thousand trips on 28 thousand to roughly 225 thousand trips on 21 thousand cards.

Overground Lines

The Overground lines as they currently exist are listed in Table 2. Once assigned to the Overground network, each trip is assigned to one of the Overground lines if its origin and destination are on the same line. Otherwise, it is labelled as a "Transfer" trip.

Line	Abbreviation	Terminals
North London Line	NLL	Stratford ↔ Richmond
Gospel Oak – Barking Line	GOB	Gospel Oak↔ Barking
Watford DC Line	WAT	Watford Junction ↔ London Euston
West London Line	WLL	Willesden Junction ↔ Clapham Junction

Table 2: The London Overground Lines

Determining Scheduled Journey Time

For a trip assumed to have used only the Overground, determining its Scheduled Journey Time given the structure of the Overground network and accompanying timetables is relatively straightforward. The approach taken is similar to the "Forward Search Algorithm" described in (Huang and Peng, 2002), but makes use of a different core path finding algorithm, namely that of (Hao and Kocur, 1992). In the context of this report it suffices to say that the algorithm employed finds the earliest possible time that the passenger could have arrived at his or her destination using only the London Overground, given the touch-in time and the assumptions listed above. It is most easily and intuitively understood through examination of Figure 2.

Software

Most of the work on this study was conducted with R, an open source programming language and environment for statistical modelling and graphics (see http://www.r-project.org for more information). Some larger queries were conducted in an Oracle database maintained within TfL for Oyster-based research. Implementation details are not the focus of this report, so the discussion will in general avoid further reference to particular software development, computational, or algorithmic strategies.

4 Exploratory Analysis

Here we present a number of exploratory analyses that become possible once the above methodology has been applied to our samples of Oyster data.

Trip Distributions by Time, Line, and O-D Pair

Figure 3 presents a number of exploratory plots, showing the distribution of Overground Oyster trips by time of day, by journey time, by line, and by origin-destination pair. These distributions show that, for example:

- Overground trip volumes follow a familiar pattern over the course of the day, with a sharp morning peak, a somewhat flatter evening peak, and a peak-to-base ratio of between 3.0 and 4.0. We also see that the evening rush hour begins at roughly 15:00, somewhat before the 16:00 official start of the PM Peak timeband.

- The median Observed Journey Time for Overground trips is close to 30 minutes, with fewer than 30% taking 20 minutes or less and close to 20% taking 40 minutes or more. Scheduled Journey Times are, in general, slightly less than Observed Journey Times, suggesting that we will find Excess Journey Time in this data.

- The Scheduled Travel Time (the time passengers would spend on board the train if the service ran perfectly to schedule) has a median of over 15 minutes, with 40% of trips spending 20 minutes or more on a vehicle.

- The North London Line serves the majority, over 60%, of London Overground-only trips. A substantial fraction of trips, roughly 14%, involve a transfer between Overground lines.

- All of the top 12 origin-destination pairs on the London Overground are in the corridor between Stratford and Camden Road. Of the top 25 O/D pairs, 19 are on the North London Line, and the other three Overground lines are each represented at least once in the remaining six.

Scheduled Waiting Time Distributions by Line and Timeband

Measuring Scheduled Waiting Time for Overground trips gives us insight into the degree to which London Overground passengers time their arrival at the station according to the published timetable. Figure 4 plots the distribution of Scheduled Waiting Time by line and by timeband. Headways on the Overground vary by line and by time of day, so trips are characterized in this plot by the ratio of Scheduled Waiting Time to headway. A low value indicates that the passenger arrived shortly before the next scheduled departure, while a high value indicates that they just missed the preceding departure. We see in these distributions that:

- Passengers on the GOB and Watford DC line tend to arrive closer to the next scheduled departure, especially earlier and later in the day. This is consistent with conventional wisdom for services with 20 to 30 minute headways.

- The distribution is much less skewed for the North London Line in the AM Peak, Inter-peak, and PM Peak timebands, during which it generally runs at 15 minute headways. This indicates that many passengers treat the North London Line as a turn-up-and-go service.

- The behaviour of passengers on the West London Line is not clear cut.

Appendix C discusses this issue in depth, including the effect that turn-up-and-go behaviour has on vehicle crowding levels and timetabling practice.

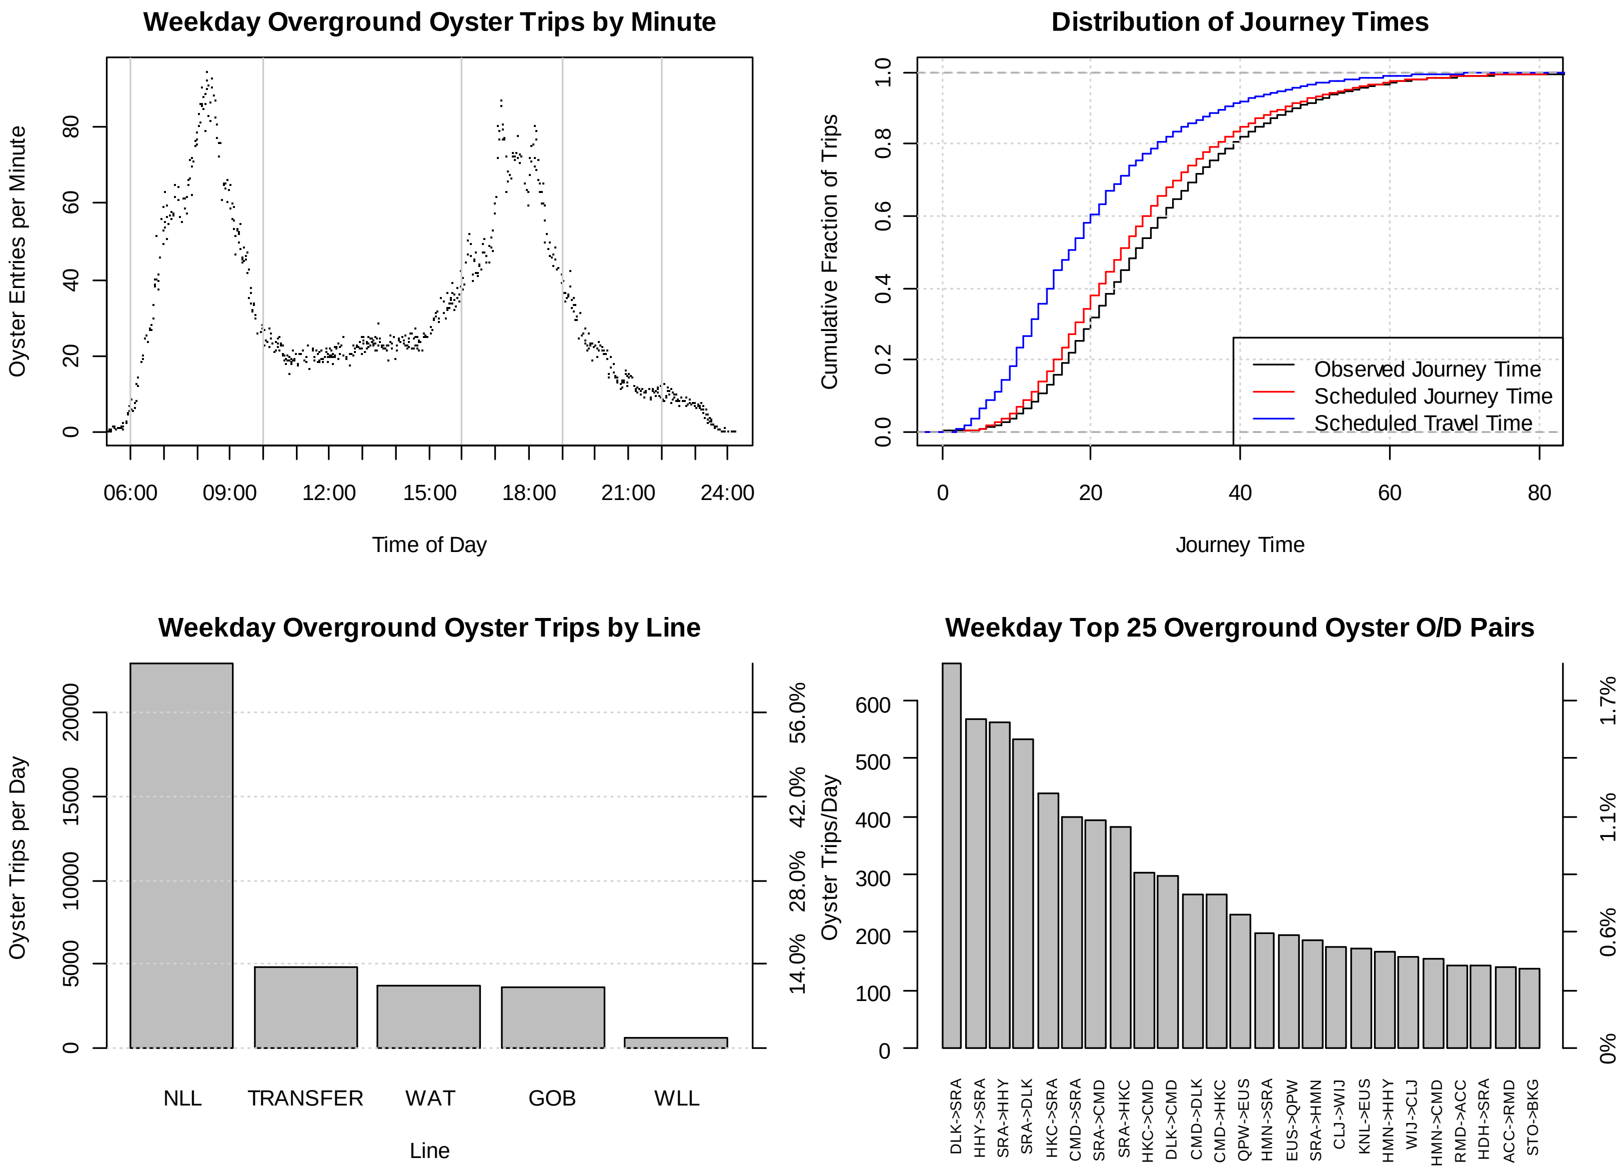

Figure 3: Exploratory Analyses of Overground Oyster Data. The four figures show the distribution of trips on the London Overground by: (in clockwise order from the top left) Time of Day, Journey Length, Line, and Origin-Destination Pair

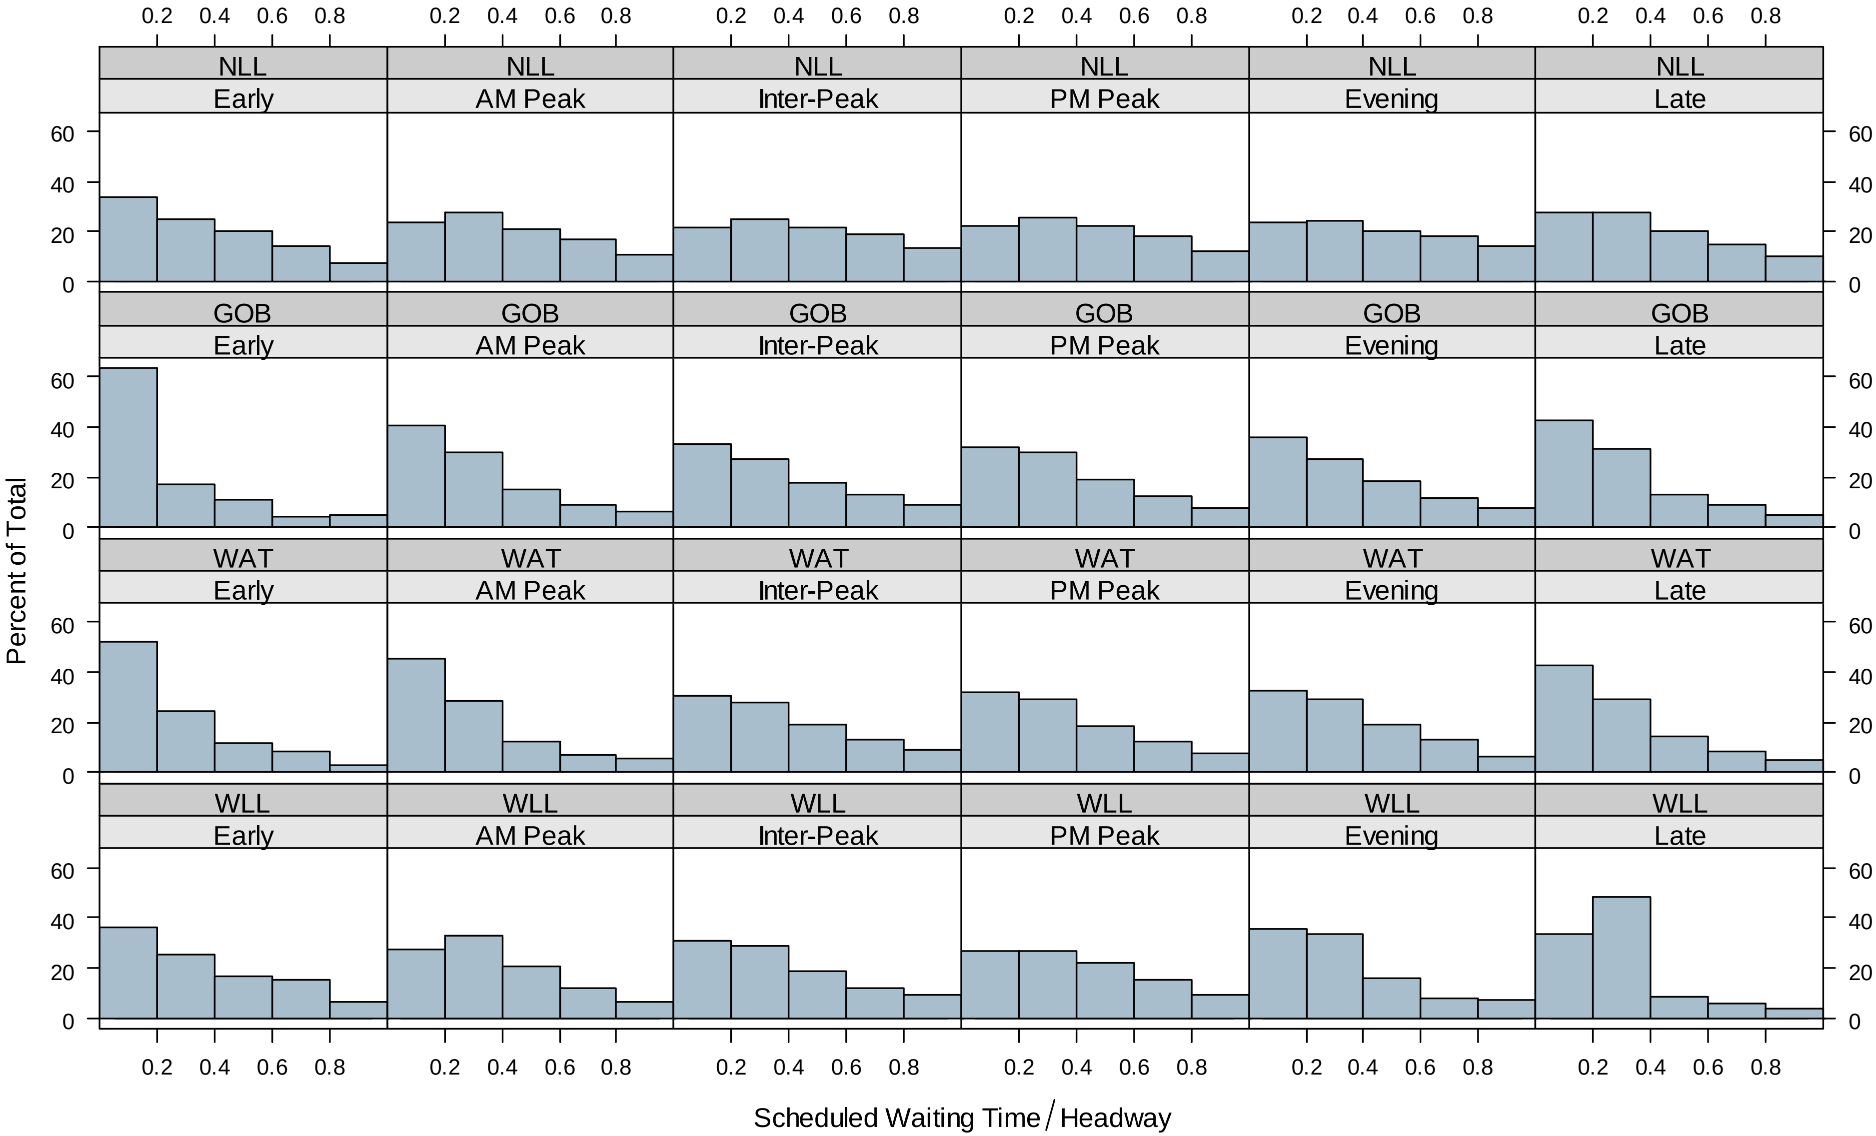

Figure 4: Distribution of Scheduled Waiting Time, normalized by headway, by Overground line and timeband

5 Validation

Once the Observed, Scheduled, and Excess Journey Times were calculated for each trip in the sample, we felt it necessary to validate these measurements before they could be aggregated to a level appropriate for use in a performance or contract management scenario. Two types of validation were performed – firstly to check that the above O/D classification scheme is reasonable, and secondly to be sure that the excess journey times as measured by Oyster accurately reflect events on the ground.

Classification of Origin-Destination Pairs

To validate the classification of O/D pairs in terms of the likelihood that trips on that pair are taken by Overground, we compared average Excess Journey Times for the different types of O/D pairs. It was found that the average EJT is negative for all types except pure London Overground (LO) journeys, broadly validating the classification scheme. Some additional in-depth analysis of specific O/D pairs was conducted, which led to a handful of adjustments to this scheme including the creation of the London Overground + London Midland Rail (LMR) type.

Effects of Operations on Passengers

Because of the complex and dynamic nature of even a single day's rail operations and passenger journeys, a graphical approach was used to validate that the Oyster-based Observed and Excess Journey Time measurements reflect real events of the sort we hope to capture. An example of the type of plot used is shown in Figure 5.

In this plot, the X axis represents time and the Y axis represents the distance travelled along the North London Line. Each blue line travelling Northwest through the plot shows the scheduled progress of one Weekday service. The size of the slanted hash marks represent the number of Oyster trips, in the Stratford → Richmond direction, that exited a given station at a given minute of the day on Thursday 3 April, 2008. The colour of each hash mark indicates the average Excess Journey Time for the trips it represents – the redder the mark, the more delay incurred.

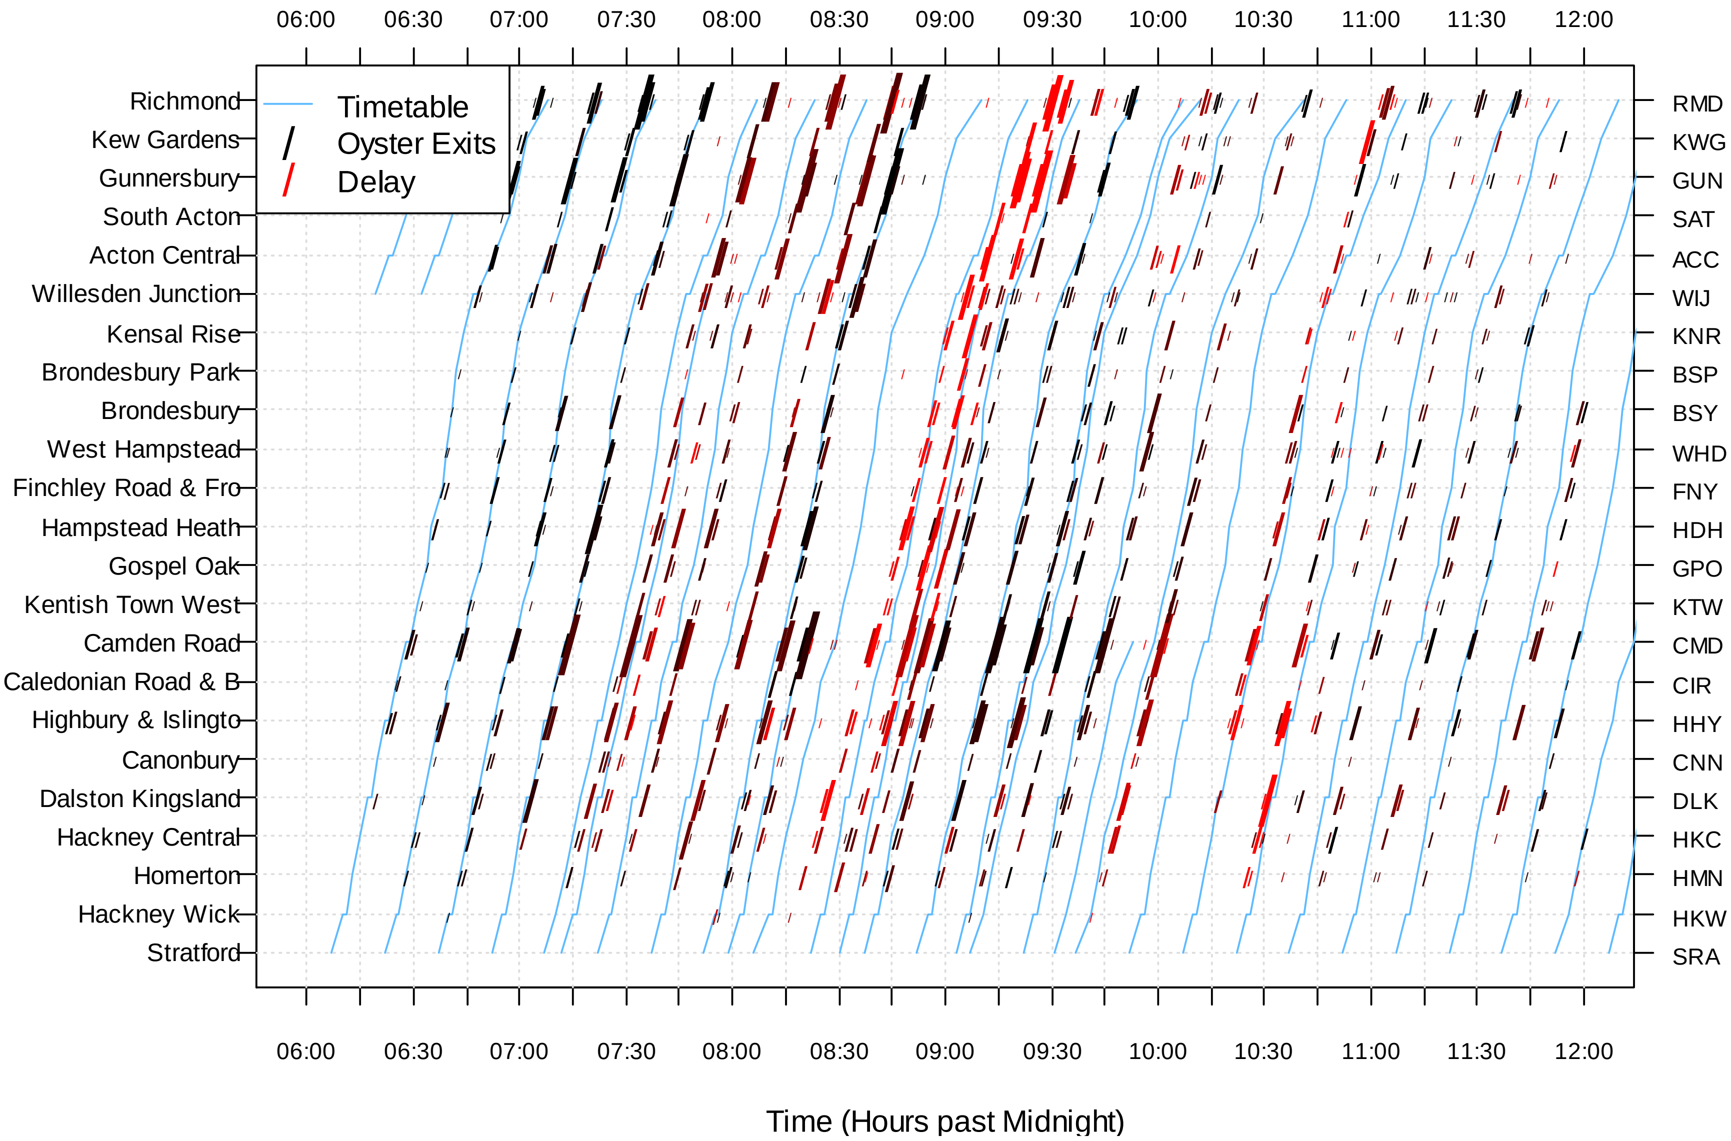

Figure 5: Time-Distance Plot for North London Line in the Stratford to Richmond direction, with Oyster exits coloured by Excess Journey Time, for the Morning of Thursday 3 April, 2008

Using Oyster data alone, there is much we can guess from this plot about the actual operation of the service on the day in question. Visually correlating trails of Oyster exits with timetabled trips, it appears that:

- Trains departing Stratford before 7:00 were relatively less patronized and ran to schedule
- Starting with the 7:07 service from Stratford there are slight delays, which become severe (i.e. bright red) for the 8:06 and 8:22, and perhaps 8:30, and 8:37 trips.
- The 8:52, 9:03, 9:22, and 9:38 services ran relatively smoothly, at least as far as Willesden Junction
- The 9:31 shuttle to Camden Road may not have run at all, as reflected by the late (i.e. red) passengers as far as Camden Road on the 9:38 service
- The 9:52 from Stratford ran extremely late or not at all
- By the 11:07 departure from Stratford, the service had approximately recovered.

Fortunately, Network Rail keeps detailed records of train delays on the network, so the above hypotheses can be judged against the true record of events. Looking in the so-called TRUST database that is used to calculate PPM, we find that:

- No North London Line trains departing Stratford before 7:00 on Thursday 3 April, 2008 were more than a minute late arriving at their destinations
- The trains departing Stratford between 7:00 and 8:00 arrived at their destinations between -2 and 9 minutes late
- The 8:06 and 8:22 trains from Stratford arrived at Richmond 18 and 10 minutes late, respectively, because of severe door problems on the 8:06 train.
- The 8:52, 9:03, 9:22, and 9:38 services were -2, 5, 2, and 0 minutes late at their destinations, respectively
- The 9:31 shuttle was in fact cancelled, because the vehicle scheduled for this trip was used instead for the 9:38 service, the original vehicle for which was delayed inbound from Richmond by a sick passenger
- The 9:52 departure from Stratford was actually short-turned on its incoming trip at Hackney Wick, a knock-on effect of the above sick passenger, and arrived 11 minutes late at Richmond
- All trains departing Stratford between 11:00 and 12:00 were between -2 and 3 minutes late at their destinations

Particularly compelling about this example is the cancellation of the 9:31 train and the short-turn of the 9:52 train. The precise effect of these control actions on passengers, relative to a right-time service, is captured in the Excess Journey Time measure, and is evidenced in the plot.

Validation of this sort was performed for each of the 10 weekday mornings from 31 March 2008 to 11 March 2008, inclusive, with similar results to the above in each case.

6 Aggregate Measures of Delay and Reliability

In this section, we present measurements of delay and reliability on the London Overground, in the form of different types of aggregations of Excess Journey Time.

Delay

Figure 6 through Figure 11 show the average Excess Journey Time on the London Overground over time, by line, and by timebands of the day. While the overall EJT clearly varies between lines and timebands, there is not a clear trend up or down over the last six months of operations. A number of observations, none particularly surprising, can be made from these figures, including:

- The North London Line has the highest Excess Journey Time of all lines on the London Overground, around 3 minutes.
- On all lines, there is significant day-to-day variability of Excess Journey Time
- For the whole London Overground, and for the North London Line in particular, Excess Journey Time is worst in the AM and PM Peak timebands.
- A slightly increasing trend can be observed in the Excess Journey Time on the North London Line in the Inter-Peak timeband.

Two peculiarities are also noticeable, both of which we feel can be attributed to problems with the methodology described above.

- The West London Line shows substantially more variation in daily Excess Journey Time than the other Overground lines. We believe this is an artefact of relatively small sample sizes on the West London Line in the 5% sample – an average of 27 trips/day as compared to over 200 trips/day on the GOB or Watford DC and over 1,200 trips/day on the North London Line. Similar plots made using the 100% sample (albeit over a shorter span of time) do not show the same variability.
- The aggregate Excess Journey Time for “transfer” trips – those using two Overground lines – is often negative. We believe this is a result of incomplete classification of origin-destination pairs for these types of trips. For example, the trip from Hackney Wick to Barking is classified as a pure London Overground trip. This would require travelling all the way back to Gospel Oak on the North London Line, waiting for a 20-30 minute service, and then travelling all the way to Barking. A more realistic route would be via Stratford and West Ham on the North London Line and then Jubilee and District underground lines.

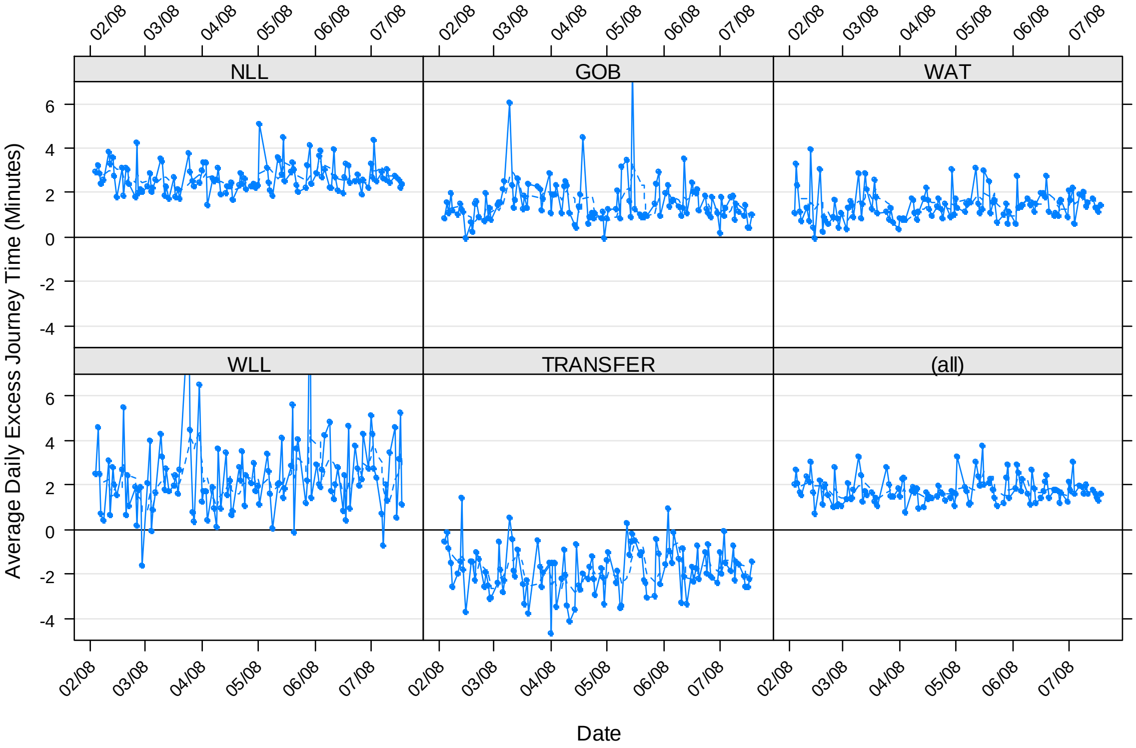

Figure 6: Daily Excess Journey Time by line, with 5-day moving average

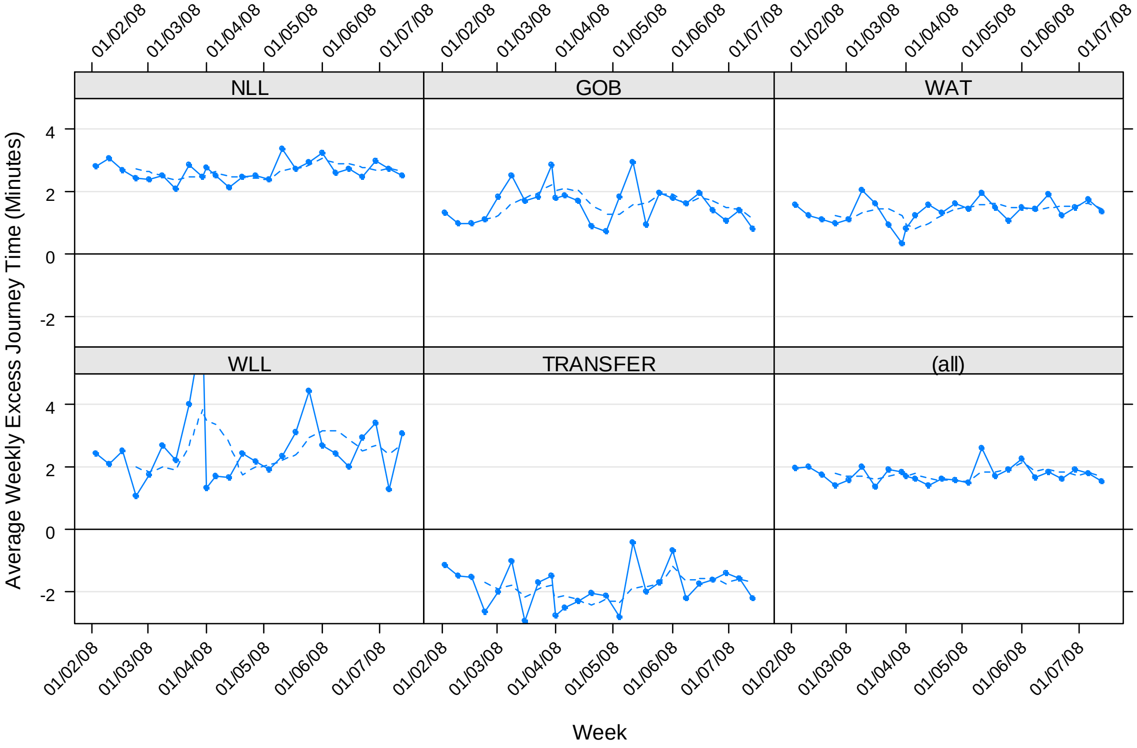

Figure 7: Weekly Excess Journey Time by line, with 4-week moving average

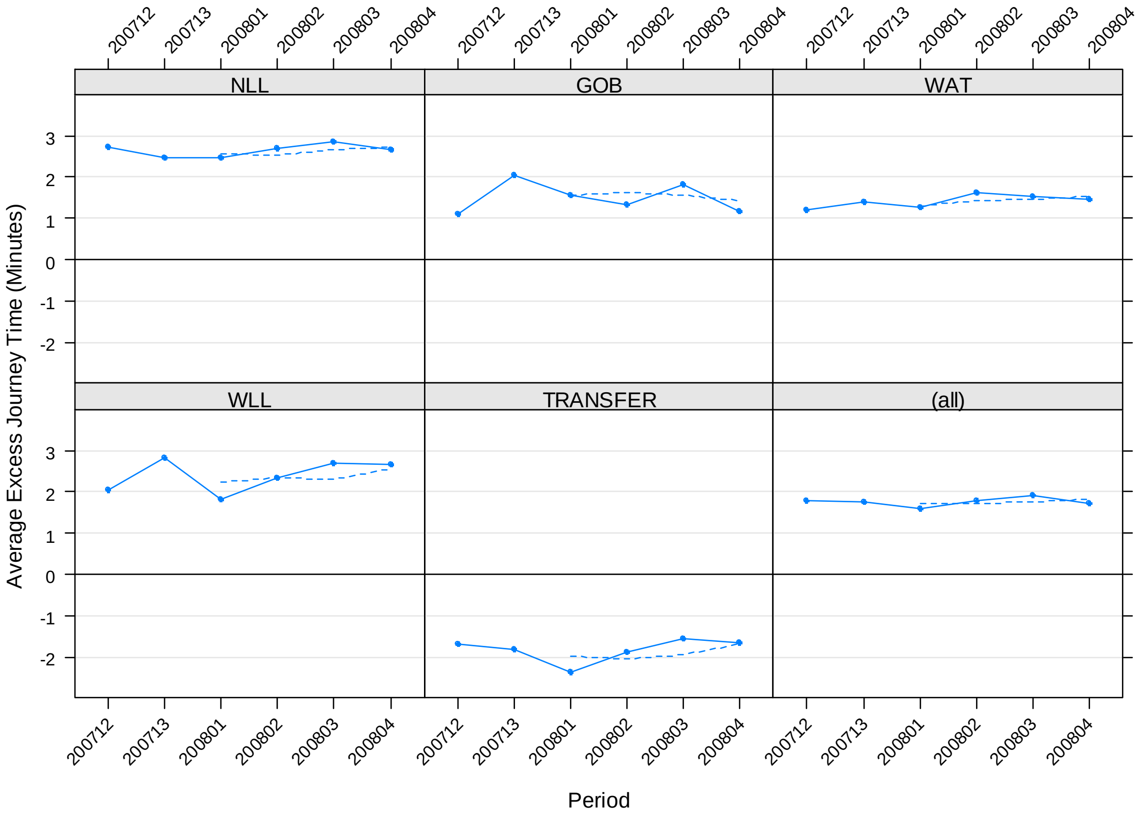

Figure 8: Period Excess Journey Time by line, with 3-period moving average

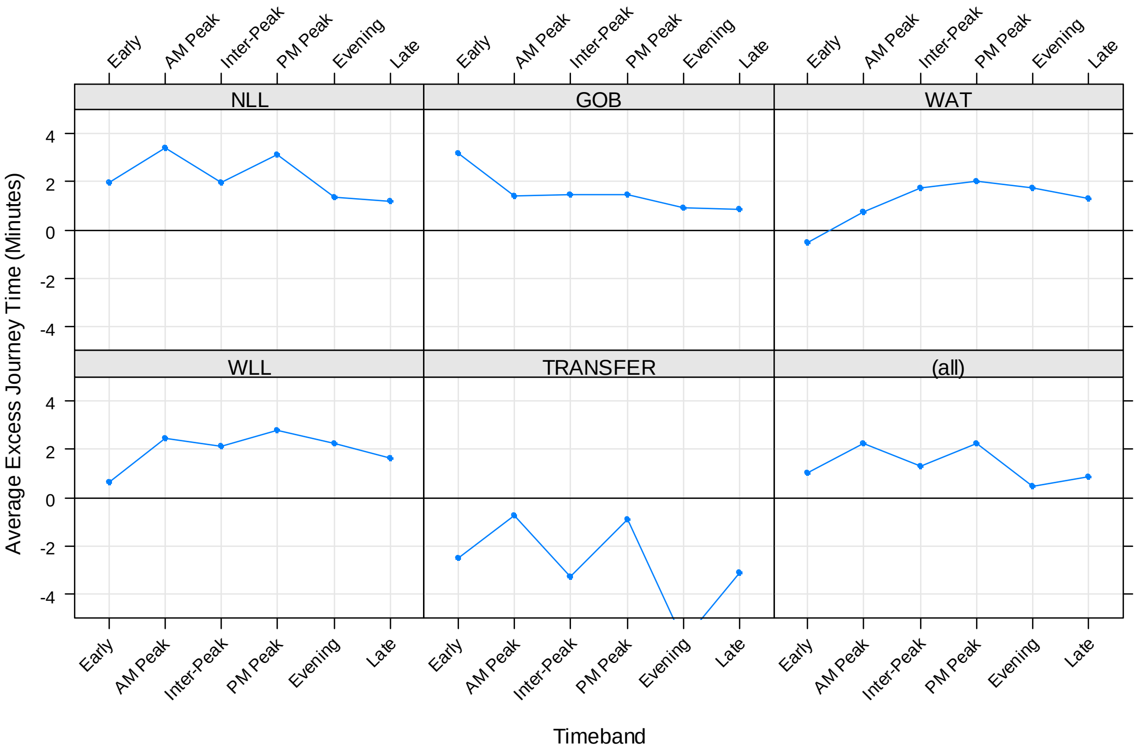

Figure 9: Timeband Excess Journey Times by Line

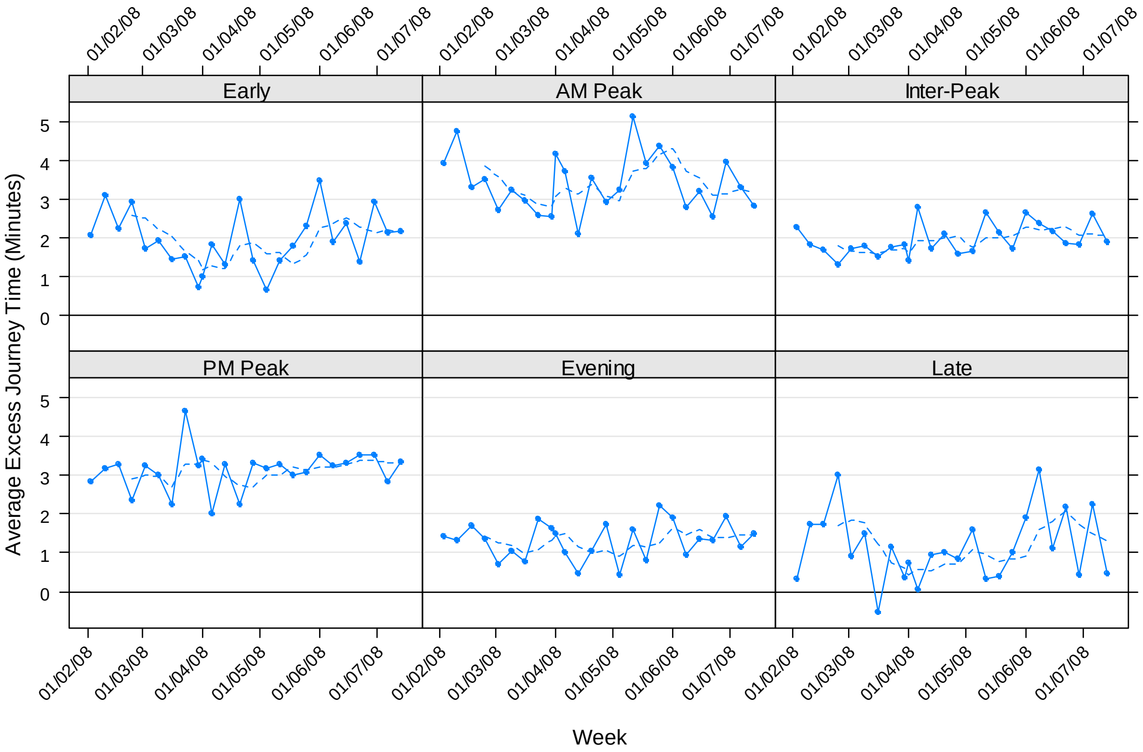

Figure 10: Weekly Excess Journey Time by timeband on the North London Line

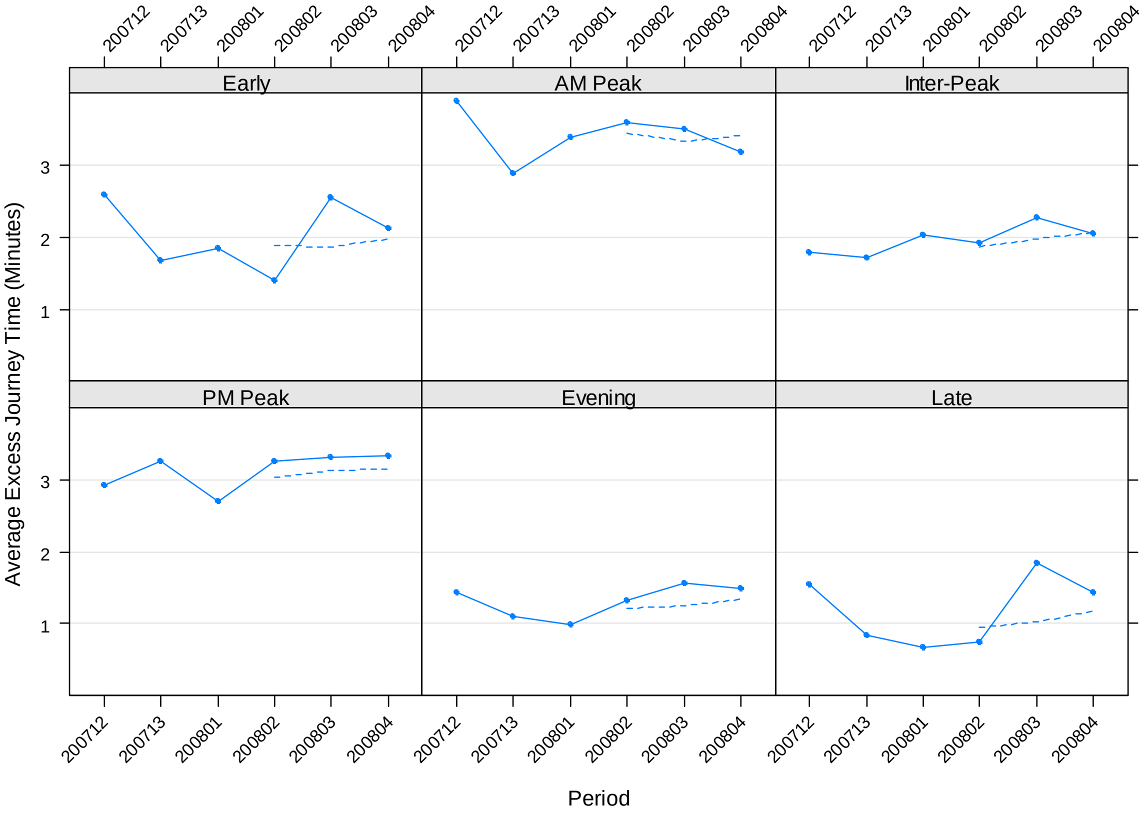

Figure 11: Period Excess Journey Time by timeband on the North London Line

Reliability

As we saw in Figure 6, even the average Excess Journey Time varies substantially from day to day. From the perspective of the passenger, this indicates a certain level of unreliability in using the Overground, in that on some days a given trip may be more delayed than on others.

We loosely define reliability as some function of the probability of a given trip being some amount longer than would be expected. Thus, measures of the distribution of Excess Journey Time can be used to indicate reliability. Figure 12 and Figure 13 show distributions of Excess Journey Time for the London Overground in the form of box plots, which indicate the 10th, 25th, 50th, 75th, and 90th percentiles of a given distribution.

To understand how these distributions indicate reliability, it is instructive to compare Figure 12 to Figure 9. For example, the average Excess Journey Time for the North London Line is about 3.5 minutes in the AM Peak but the 90th percentile is 10 minutes. Thus, customers travelling on the North London Line in the AM peak who want to arrive at their destinations on time 90% of the time must budget an additional 10 minutes beyond the published timetable, even though *on average* they will be only 3.5 minutes late. On the other hand, passengers on the GOB during the Inter-Peak period can expect to be about 1.5 minutes late, and 90% of the time will be at most 3 minutes late.

Reliability, as measured by distributions of Excess Journey Time, can and should play an important role in the evaluation of a railway's operations because it has profound impacts on the day to day experience of passengers.

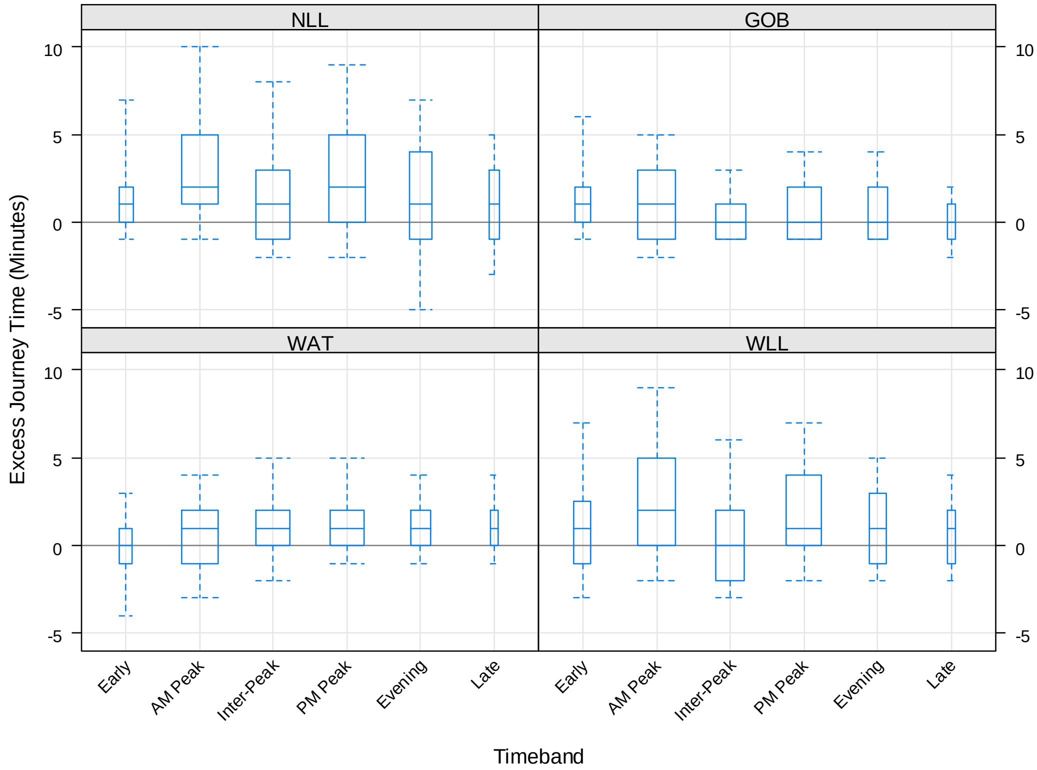

Figure 12: Distribution of Excess Journey Time over full sample, by timeband by line. The box-and-whisker plots represent the 10th, 25th, 50th, 75th, and 90th percentiles of the respective distributions.

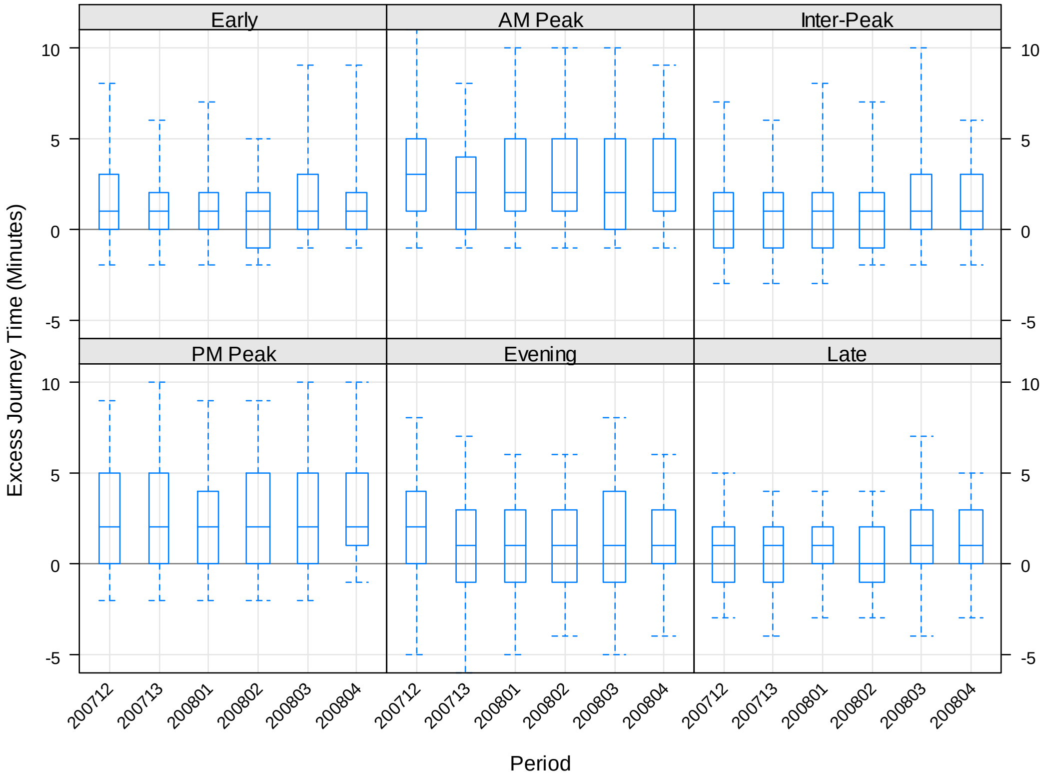

Figure 13: Period distribution of Excess Journey Time on the North London Line by timeband.

7 Relationship of Excess Journey Time to Public Performance Measures

This section compares Excess Journey Time results to PPM measurements, first graphically and then through linear regressions. On the one hand, we expect to see a substantial level of correlation. On the other hand, perfect correlation would indicate that no extra information is captured by the Excess Journey Time measurements.

Graphical Comparison

Figure 14 and Figure 15 plot Excess Journey Time and PPM over the whole London Overground network for the longest period of time in our sample. We see that over time they tend to move similarly and that their relationship is somewhat linear.

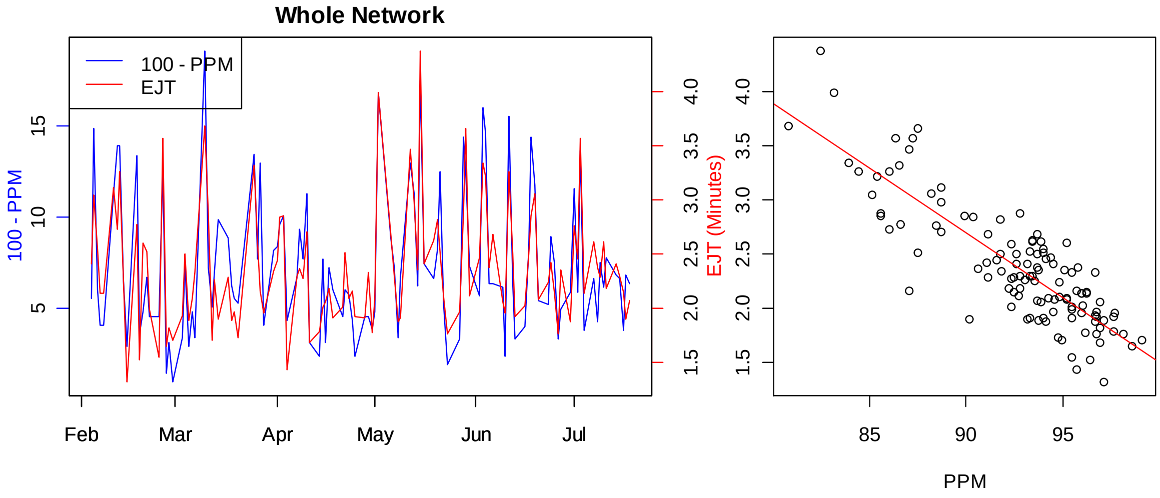

Figure 14: Daily EJT vs. PPM for the London Overground Network

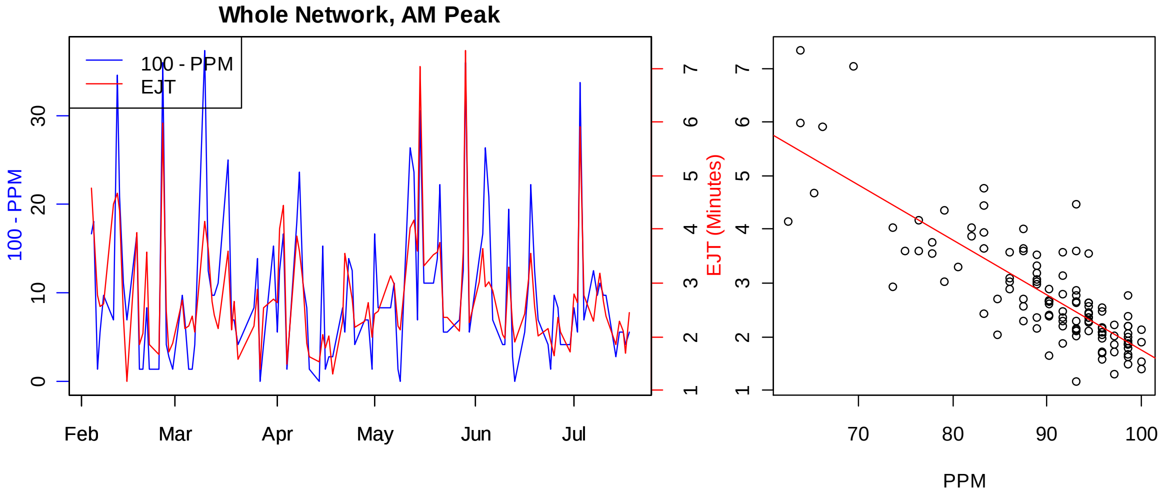

Figure 15: Daily EJT vs. PPM for the London Overground Network in the AM Peak

However, there is clear variation of EJT around the linear fit of PPM, indicating that there is additional information captured in Excess Journey Time measurements. This variation tends to increase as PPM decreases, particularly in the AM peak, which suggests that PPM is a less effective measure of passenger experience as conditions worsen.

Linear Regression Models

A number of linear models were estimated to test the relationship between aggregate Excess Journey Time and the indicators provided by the PPM/TRUST system. One hope was that if a sufficiently accurate model were found, it could be used with confidence to project EJT figures into the past based on archived PPM data. Unsurprisingly, we were not so lucky.

The first set of models, shown in Table 3, were estimated for the whole Overground for the entire day and for the AM Peak. These simple models test EJT as a linear function of the two primary outputs of PPM:

- PPM Failures – the number of scheduled services that did not reach their destination less than 5 minutes late on a given day. This number was used rather than the PPM percentage itself so that the estimated coefficient would indicate the effect on EJT of each problem train.
- Total Delay Minutes – the total number of vehicle delay minutes accumulated by services that ran on a given day.

The simplest possible model included only an intercept term and PPM Failures. Somewhat more explanatory were models with PPM Failures and Total Delay Minutes. In all four models all coefficients were of the expected sign (positive) and statistically significant, but none explained more than 72% of the variation in EJT.

	Whole Overground Network, All Day				Whole Overground Network, AM Peak			
	PPM Only		PPM + Delay Min.		PPM Only		PPM + Delay Min.	
Coefficient	Estimate	Pr(>|t|)	Estimate	Pr(>|t|)	Estimate	Pr(>|t|)	Estimate	Pr(>|t|)
(Intercept)	0.881	0.000	0.742	0.000	1.184	0.000	0.977	0.000
PPM Failures	0.029	0.000	0.018	0.000	0.145	0.000	0.095	0.000
Total Delay Minutes			0.001	0.012			0.005	0.001
Adjusted R^2	0.706		0.719		0.675		0.701	

Table 3: Estimation results of linear models for EJT on the whole Overground network

We estimated the PPM + Delay Minutes model as well as a more complicated model for the North London Line in the AM and PM Peaks, with the results shown in Table 4. The simpler model is more explanatory for the North London Line that it was for the whole Overground Network, which stands to reason because it was in this case estimated over a less heterogeneous set of observations.

	North London Line, AM Peak				North London Line, PM Peak			
	PPM + Delay Min.		Complex Model		PPM + Delay Min.		Complex Model	
Coefficient	Estimate	Pr(>|t|)	Estimate	Pr(>|t|)	Estimate	Pr(>|t|)	Estimate	Pr(>|t|)
(Intercept)	1.832	0.000	1.243	0.000	1.752	0.000	1.753	0.000
PPM Failures	0.142	0.000			0.068	0.002		
Trains 1 - 4 min late		0.000	0.090	0.000			0.012	0.478
Trains 5 - 9 min late			0.119	0.000			0.077	0.011
Trains 10 - 19 min late			-0.026	0.765			0.097	0.026
Trains 20 + min late			0.161	0.396			0.305	0.420
Trains terminated before destination			0.141	0.089			0.095	0.298
Trains started ahead of origin			0.071	0.479			0.119	0.145
Full cancellations			0.662	0.000			0.251	0.033
Total Delay Minutes	0.012	0.000	0.011	0.000	0.014	0.000	0.012	0.000
Adjusted R^2	0.773		0.843		0.828		0.827	

Table 4: Estimation results of linear models for EJT on the North London Line in the AM and PM Peak

The more complex models include trains that were less than 5 minutes late, and break PPM Failures apart by all possible reasons for failing the PPM, the natures of which should be self explanatory from the table. For the PM Peak the explanatory power does not increase between the simple and complex model, but the magnitude and signs of coefficients are in line with expectations. For trains late by different numbers of minutes, the coefficient increases with the amount of lateness, though is not significant at below the 40% level for trains 1-4 minutes or 20+ minutes late. Trains started ahead of their planned origin are somewhat worse than trains terminated before their destination, but these factors are only significant at the 30% and 14% level, respectively. Full cancellations have the most effect (per train) on EJT of all of the terms that are highly significant.

For the AM peak we see a substantial increase in explanatory power, with the adjusted R^2 increasing from 0.77 to 0.84, and a similar pattern in the coefficients as in the PM peak. In this case, the number of trains fully cancelled has a quite high coefficient and is significant at below the 1% level. The number of trains terminated before their destinations has more effect than trains started ahead of their origins, and is significant at less than the 10% level. However, the model is not entirely satisfying in that the number of trains 10-19 minutes late has a negative coefficient (though significant only at the 76% level), and that the coefficients for trains 20+ minutes late and for trains started ahead of their origins are not significant below the 45% level.

These were as good models as we were able to get with the data in hand, but it is not felt that any of them are accurate enough to meaningfully project EJT in the past (or future) from PPM and TRUST data.

8 Ambiguity and Opportunity

The creation of an integrated transport network is among the top priorities for transport investment and planning in London. We have been quite successful in this endeavour in the Overground case, as indicated by the fact that roughly half of all Oyster trips originating at stations served exclusively by the Overground terminate at stations with no Overground services at all (and vice versa).

At the same time, much of the focus on performance measurement and management is directed at individual modes or components of the network. This report so far is no exception to that rule. Despite the preliminary results, which we find quite encouraging, the use of the Oyster system as the primary source of performance measurement poses some problems.

As mentioned above, Oyster data record only the locations and times of entry and exit for each trip. Thus it fully captures information about the end-to-end journey times experienced by passengers, but is inherently ambiguous in the following respects:

- Journey time on each individual leg of any trip involving the use of more than one service
- Choice of route or path for any trip with a number of likely possible routings between origin and destination
- Choice of service on interavailable routes
- Time spent waiting on platforms as distinct from time spent travelling on trains
- Time spent in access, egress, or interchange between services

These are serious obstacles to accurately and robustly measuring the performance of any one specific piece of the overall network on which Oyster is used, including the London Overground. However, they each reflect distinct opportunities in the broader context of measuring the passenger experience *across* the whole of London’s public transport network.

This section discusses the above ambiguities, both in terms of the problems they present for measuring the performance of the Overground and in terms of the opportunities they present for measuring the integrated performance of the entire Oyster-enabled portion of London’s transport system. This is a particularly salient moment to begin to consider such a prospect, as Oyster is in the process of being expanded to the National Rail network in the Greater London Area.

Transfers

Figure 3 shows that 14% of the trips on the Overground network in our samples require the use of more than one Overground line. Any Excess Journey Time of these trips is not easily attributable to the performance of any one line, let alone a particular incident, but at least is attributable to the Overground in general.

The situation is much less straightforward for trips using the Overground and one or more other Oyster-enabled rail modes, as illustrated in Figure 16. These trips constitute a substantial portion of total traffic on the Overground, as it is estimated that roughly half of all Oyster trips originating at stations served exclusively by the Overground terminate at stations with no Overground services at all.

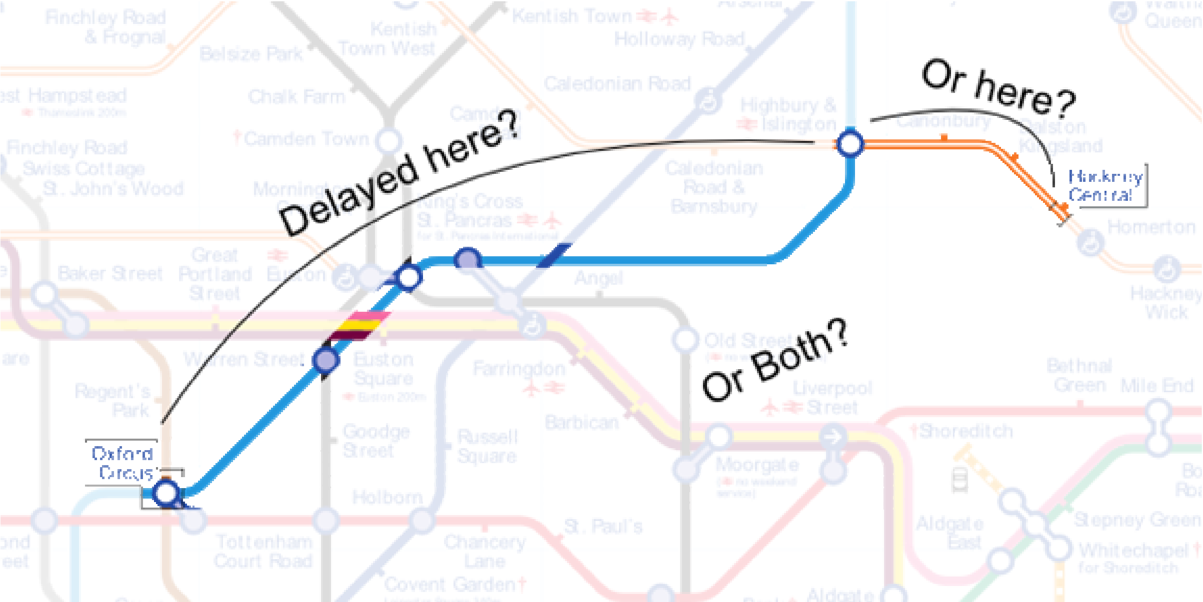

Figure 16: Illustration of an Overground trip with transfer to another rail network

Advanced statistical inference techniques, discussed further in the section on Future Research, could be used to estimate the delay incurred on passengers by different parts of the network, but the true value may not lie in that course of action.

The opportunity presented by Oyster data on these kinds of trips is the ability to understand the experience of passengers using multiple parts of the network in their journeys. This information is not provided by any of the mode-specific performance measures.

Of course, delays in transfer trips will, at some level, need to be attributed to problems in a specific part or parts of the network, so it is likely that inference and mode-specific performance regimes would still be of value.

Path Choice

When a specific Oyster trip has only one likely path between the origin and destination, it can be straightforward to set an expectation of journey time. However, when multiple likely paths exist, it is not always so easy. A given trip will have a different Excess Journey Time depending on which path it is measured against.

Consider the example shown in Figure 17 – a trip from Hackney Central on the North London Line to Tottenham Court Road. If the expected journey time is 30 minutes via Highbury & Islington and Oxford Circus but 36 minutes via Stratford with fewer transfers, then the Excess Journey Time is quite different depending on which path was taken. The difference in EJT will always be 6 minutes, and if the observed time is between the two expected times (e.g. 32 minutes) it will be positive for the former path (+2) and negative (-4) for the latter.

Figure 17: Illustration of an Overground trip with path choice

Since Excess Journey is averaged, or possibly summed, in the aggregation process, this problem should be quite straightforward to cope with as long as we have a probabilistic model of path choice. Probabilities for each path, of the sort developed for Pay As You Go revenue allocation the Oyster Extension to National Rail project, would allow the aggregate sums or averages to be weighted appropriately.

In the broader measurement context, multiple reasonable paths make travel on a given origin-destination pair more robust over time to problems in the network. This means that Excess Journey Time measurements for the whole Oyster network should help focus attention on the experience of passengers who are captive to a single route for their travel.

Interavailable Services

There are several places where the Overground runs in parallel to the London Underground or National Rail services, stopping at the same stations. These so-called “interavailable services” are in some sense just a special case of path choice, and could be given the same treatment.

Waiting vs. In-Vehicle Time

In assessing the passenger experience it is often important to understand how much time was spent waiting versus riding in vehicles, usually for the sake of weighting these times differently. Oyster data does not explicitly offer this information, so we face a similar problem was described above regarding trips with legs on different networks, and similar possible solutions.

Access/Egress/Interchange

Access, egress, and interchange times (excluding additional waiting time) are not typically affected by the operations of a given service. They are subject to pedestrian congestion and crowding effects and often have significant impact on decisions regarding investment in station facilities. As with waiting time, they are not directly measurable with Oyster data.

In terms of total journey times, however, access and interchange times do not in and of themselves represent a problem unless they cause a passenger to miss the train he or she would have otherwise caught. Thus, measuring end-to-end journey times should more accurately capture the negative effects of crowded stations or broken escalators on real passenger journeys than do direct measurements of access and interchange times.

9 Prospects for Application

The work described here has a number of potential applications in near- and mid-term timeframes.

Reporting of London Overground Performance Consistent with TfL Objectives and Other TfL Modes

The first priority set out in TfL's current Business Plan is to "Improve door-to-door journey times and reliability across our transport system" (TfL 2007). Consistent with this objective, TfL's two premier public transportation modes – Underground and Buses – currently report their performance with metrics measured in units of passenger time. The London Overground on the other hand reports performance in terms of punctuality of scheduled trips, an indirect measure of the passenger experience.

The methodology developed in this report can be used to cheaply and easily measure and report the service quality of the London Overground. In addition to average delay to passengers, reliability can be quantified by measuring the distribution of delays. These measures serve not only to inform operations planning and management, but also to demonstrate to the travelling public that investment in TfL is having the desired effects.

Improved Collaborative Management of Operations on the London Overground

Sections 4 through 6 and Appendix C of this report present broad and deep data-driven analyses of the effect of the operational performance of the London Overground on the experience of its passengers. These results offer certain perspectives that are not provided by the PPM, and should be of particular interest to both the operator, LOROL, and the owner, TfL.

Regular and thorough of analysis of Oyster data should be undertaken by TfL and LOROL in their shared roles in the management and improvement of the London Overground. Together they can more effectively assess problems and find solutions. Appendix C presents one example of the usefulness of Oyster data, wherein timetabling changes are recommended to ease unbalanced passenger loads based on the observed behaviour of passenger arrivals relative to published timetables.

Measuring Performance of National Rail Services in the Greater London Area

The same methodology applied here to the Overground can be directly applied to those National Rail services that accept Oyster for payment. Acceptance of Oyster on National Rail is currently quite limited (NR 2008b) but over the next several years is to be rolled out to all National Rail services in the Greater London Area.

The data and methods described here present TfL and the Mayor of London with an opportunity to play a stronger role in representing the needs of people travelling within the Greater London Area. As the owner/operator of the Oyster ticketing system, and the political body most responsive to London constituencies, TfL and the Mayor can use Oyster-based performance measures to influence decision-making around the provision of National Rail services. They will be able to identify and help solve problems in the services provided by the Train Operating Companies, and guide concessioning and investment decisions made by Network Rail and the national government.

10 Future Research

This work suggests a number of directions for future research. Some are direct extensions, while others are tangentially related but potentially useful in their own right.

Performance Measurement Across the London's Entire Rail Network

Each rail mode or sub-network in London (e.g. Underground, Overground, DLR, the TOC's) has its own performance monitoring scheme but it is not clear if anyone is watching out for the performance of the network as an integrated whole. As indicated in Section 8 of this report, we would very much like to see Oyster-based measurement of performance from the passenger perspective applied to the wider network of rail services in London.

Research should be conducted into the applicability of Oyster-based performance metrics to multimodal journeys and networks. Some of the open questions were discussed in Section 8 of this report, but others remain. For example, on a high frequency network such as the Underground, it is likely inappropriate to compare observed journey times to published (or unpublished) timetables. We look forward to seeing the results of measures such as Excess Journey Time for trips that have multiple transfers between sub-networks and/or multiple routing options.

Statistical Inference of Trip Component Journey Times

We believe that statistical techniques can be used to infer some of the trip characteristics that raw Oyster data does not provide. Simultaneous consideration of many trips overlapping in time and space should allow us to guess such quantities as the waiting times or link-by-link travel times experienced by passengers. Potentially useful inference methods include linear regressions, quantile regressions, and maximum likelihood estimation.

Application of Path Choice Models to Performance Measurement

In general, path choices made by passengers have profound impacts on the performance of the system and on their own experiences (Guo 2008). As established in Section 8 of this report, the path chosen for a given trip can affect the correct value of Excess Journey Time, when the actual journey time is measured by Oyster with no notion of which path was actually taken. Thus, careful study is merited to see how different path choices affect measurements of delay to passengers and reliability, and how the use of probabilistic path choice models can help improve the accuracy of these measurements.

Combination with Other Data Sources

Of course, Oyster data is not the only source of data on passenger experiences. In particular, data from train signalling systems should help resolve some of the ambiguities inherent in Oyster data, and help describe passenger experiences in ways we have not considered here. Comparison of Oyster transactions to actual train positions should yield estimates of waiting and travel times for each portion of a trip. In aggregate, this approach provides estimates of actual train loads, and has been used by the BART in San Francisco for upwards of 20 years (Buneman, 1984).

Glossary of Abbreviations

***EJT* –** Excess Journey Time, a measure of journey times in excess of expected values.

***EWT* –** Excess Waiting Time, one of London Buses' primary metrics of service quality, measures passenger waiting times in excess of expected values.

GLA – The Greater London Area, the political jurisdiction of the Mayor of London and his subsidiary, Transport for London.

GOB – Gospel Oak to Barking line, one of the London Overground Services

***JTM* –** The Journey Time Metric, one of the London Underground's primary metrics of service quality, measures passenger journey times in excess of expected values.

***LB (LBSL)* –** London Buses (Services Limited)

LMR – London Midland Rail, one of the concessioned operators of National Rail services.

LO – London Overground

***LOROL* –** London Overground Rail Operations Limited, the Train Operating Company that operates the London Overground services under concession from TfL.

***LU (LUL)* –** London Underground (Limited)

***MIT* –** Massachusetts Institute of Technology, a partner with TfL in collaborative research of applications of Oyster data

MOIRA – A timetable-based model of rail market and passenger train choice on National Rail services (ORR 2008a).

NLL – North London Line, one of the London Overground Services

NR – National Rail, or Network Rail.

***OJT* –** Observed Journey Time (in this work, as measured by the Oyster smartcard ticketing system).

***PIXC* –** Passengers in eXcess of Capacity, a measure of crowding on London Overground and National Rail services.

PPM – The Public Performance Measure is the primary measure of service quality used by the London Overground and by National Rail services.

SJT – Scheduled Journey Time (in this work, as determined by the published timetable).

***SWT* –** Scheduled Waiting Time, given an arrival at a station at a given time, the amount of time the passenger would have to wait if the service ran perfectly to schedule.

TfL – Transport for London

TOC – Train Operating Company

TRUST – A database of delays on the National Rail network, with causes and attribution to responsible parties (ORR 2008b).

***WAT* –** Watford DC line, one of the London Overground Services

WLL – West London Line, one of the London Overground Services

References

(Buneman, 1984)
Buneman, K, Automated and Passenger-Based Transit Performance Measures. Transportation Research Record 992, pp 23-28. 1984, Washington DC.

(Chan, 2007)
Chan, J, Rail transit OD Estimation and Journey Time Reliability Metrics Using Automated Fare Data. Master's thesis, Massachusetts Institute of Technology, 2007.
http://dspace.mit.edu/handle/1721.1/38955

(Guo 2008).
Guo, Z., Transfers and Path Choice in Urban Public transport Systems. PhD Dissertation, Massachusetts Institute of Technology, 2008.

(Hao and Kocur, 1992)
The Hao-Kocur shortest path algorithm in described in:
http://ocw.mit.edu/NR/rdonlyres/Civil-and-Environmental-Engineering/1-00Spring-2005/BB9F54EE-0867-4714-BCF0-7FF5C4D41131/0/lec29.pdf

(Huang and Peng, 2002)
Huang, R and Peng, ZR, Schedule-Based Path-Finding Algorithms For Transit Trip-Planning Systems. Transportation Research Record 1783, pp 142-148. 2002, Washington DC.

(LB 2002)
High Frequency Calculations Explained, London Buses, 2002.
http://source.tfl/static/surface/performance/faqs/high-frequency-calculations-explained.pdf

(LB 2008)
London Buses Quality of Service Indicators: First Quarter 2008/2009, London Buses.
http://source.tfl/pdfs/current-quarter-results.pdf

(LU 2008)
Journey Time Metric page on the London Underground corporate intranet:
http://intranet.lul.co.uk/ssd/time/modelling_and_performance/jtm/default.asp

(NR 2008a)
Performance and Financial Results page on the Network Rail web site:
http://www.networkrail.co.uk/aspx/742.aspx

(NR 2008b)
Oyster Cards and National Rail page on the National Rail web site.
http://www.nationalrail.co.uk/times_fares/london/oystercard.html

(ORR 2008a)
MOIRA page on the Office of Rail Regulation web site:
http://www.rail-reg.gov.uk/server/show/ConWebDoc.8015

(ORR 2008b)
TRUST page on the Office of Rail Regulation web site:
http://www.rail-reg.gov.uk/server/show/ConWebDoc.8093

(TfL 2007)
TfL Business Plan 2005/06 to 2009/10, December 2007
http://www.tfl.gov.uk/assets/downloads/corporate/businessplandec2007.pdf

Appendix A: Sample London Overground Period PPM Report

Part of what you get in a period PPM report is:

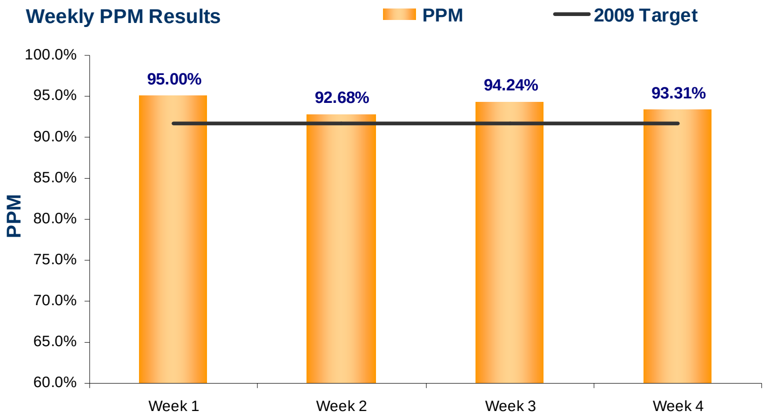

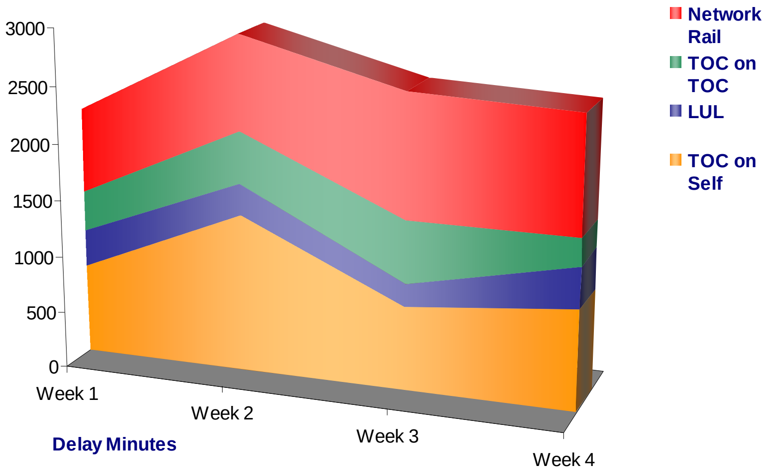

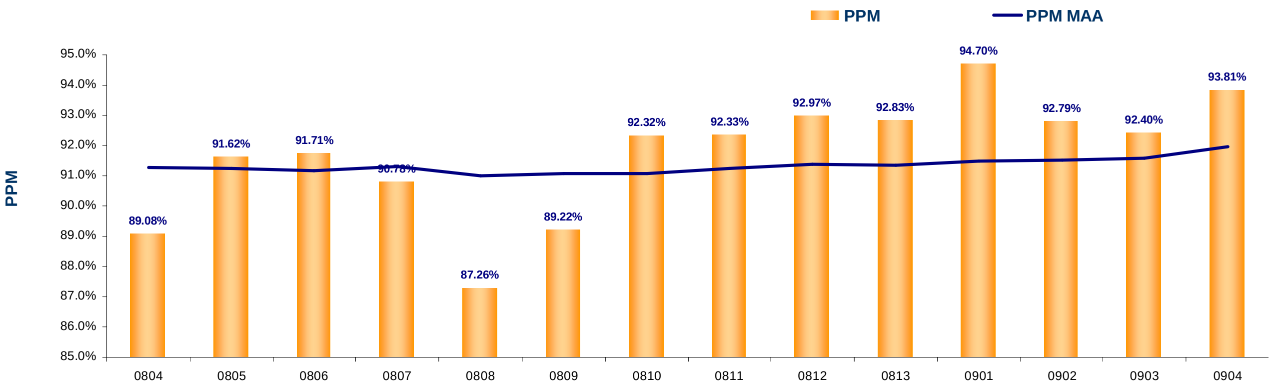

Appendix B: Sample Oyster Data

			Raw Data								Processed Fields			
Card ID	Sequence Number	DayKey	Origin	Destination	Tap In Time	Tap Out Time	Subsystem ID		Origin Code	Destination Code	Timeband	Date	Tap In Time	Tap Out Time
10001662	6403	10317	1365	1693	399	434	3		RMD	GPO	Early	2008-03-31	06:39	07:14
10001662	6405	10317	1693	1365	959	1000	3		GPO	RMD	Inter-Peak	2008-03-31	15:59	16:40
10004545	5330	10317	1690	1349	661	704	3		BSY	SRA	Inter-Peak	2008-03-31	11:01	11:44
10004545	5332	10317	1349	1690	1205	1249	3		SRA	BSY	Evening	2008-03-31	20:05	20:49
98064	97	10317	94	1707	566	599	263		HHY	BSP	AM Peak	2008-03-31	09:26	09:59
10011239	10539	10317	94	1692	422	441	263		HHY	HDH	AM Peak	2008-03-31	07:02	07:21
98335	295	10317	1695	1697	480	502	263		HKC	CMD	AM Peak	2008-03-31	08:00	08:22
98335	297	10317	1697	1695	1031	1052	263		CMD	HKC	PM Peak	2008-03-31	17:11	17:32
98441	213	10317	1698	1696	544	552	263		DLK	HMN	AM Peak	2008-03-31	09:04	09:12
98509	423	10317	1710	1689	436	459	263		FNY	ACC	AM Peak	2008-03-31	07:16	07:39
98509	425	10317	1689	1710	1024	1052	263		ACC	FNY	PM Peak	2008-03-31	17:04	17:32
98594	334	10317	1697	94	519	535	0		CMD	HHY	AM Peak	2008-03-31	08:39	08:55
98607	570	10317	1689	1698	436	479	263		ACC	DLK	AM Peak	2008-03-31	07:16	07:59
98607	572	10317	1698	1689	930	970	263		DLK	ACC	Inter-Peak	2008-03-31	15:30	16:10
10000520	4538	10317	1365	1688	973	1009	263		RMD	WIJ	PM Peak	2008-03-31	16:13	16:49
31938	26	10317	1689	1365	471	487	3		ACC	RMD	AM Peak	2008-03-31	07:51	08:07
10024007	6406	10317	1365	1343	400	409	3		RMD	GUN	Early	2008-03-31	06:40	06:49
10024007	6408	10317	1343	1365	913	922	3		GUN	RMD	Inter-Peak	2008-03-31	15:13	15:22
10036267	1964	10317	1688	1691	969	990	20		WIJ	WHD	PM Peak	2008-03-31	16:09	16:30
10015941	53365	10317	94	1696	469	485	263		HHY	HMN	AM Peak	2008-03-31	07:49	08:05

Table 5: Sample Oyster Data

Table 5: A small sample of selected columns Oyster data. To the left of the grey column are original values from the Oyster sample. To the right are columns added for human readability. The Card ID column is an anonymized version of the original card identifier.

Appendix C: Random Arrivals, Even Headways, and Vehicle Crowding

We saw in Figure 4 that passengers on the North London Line arrive near-uniformly over any given interval between consecutive services. The sort of behaviour, known as "random arrivals," has direct implications for timetabling.

Well-established transport theory establishes that, under random arrivals, the expected platform waiting time for any given passenger is minimized by provision of even headways. Intuition also tells us that, when passengers arrive at a uniform rate over time, the longer the leading headway for a given service, the more passengers we can expect to board. Thus, even headways have the effect of balancing passenger loads on successive services as well as minimizing passenger waiting time.

As shown in Figure 18, the North London Line is throughout the morning a 15 minute service. However, during the AM peak a number of additional trains, departing from Stratford at 7:12, 7:59, 8:30, and 9:07, and 9:37, are scheduled between the regular services,. These so-called "PIXC Buster" trains are intended to relieve crowding during the busiest times of day on the railway. They were slotted this way between the regular services when Transport for London offered to pay Silverlink, the previous operator of this railway, to run them.

Figure 19 shows that these services do achieve their desired effect of lessening crowding on some services. It also shows that the resulting loads are far from balanced. The effect is somewhat masked by the changing levels of demand throughout the AM Peak, but is clearly seen in the series of trains departing Stratford between 7:52 and 9:52, inclusive.

The clear pattern is as follows:

- A service with a leading headway of 15 or 16 minutes with a relatively high number of boardings
- Two successive services of headways less than 15 minutes, each with substantially fewer boardings
- A 15 minute service with, again, a relatively high number of boardings.

This analysis suggests that the peak period timetabling of the North London Line should be re-examined, and that an emphasis should be placed on even headways between scheduled services. It is likely that even headways of 9 or 10 minutes, rather than 15 minutes followed by 7 minutes and 7 minutes, would result in substantially more balanced passenger loads.

For reference, of the passengers in our sample boarding the North London Line Westbound in the AM Peak between Stratford and Caledonian Road, inclusive, 45% alight beyond Camden Road. Note that in this analysis, passengers were assigned to board the first service they should see according to the timetable, even if it is not a through service (i.e. the 8:06 and 9:03 shuttles to Camden Road and the 7:12 and 8:30 specials to Clapham Junction via Willesden Junction). Under the assumption that passengers would wait for the first service that calls at their final destination, the loads shown in Figure 19 would only be *more* unbalanced.

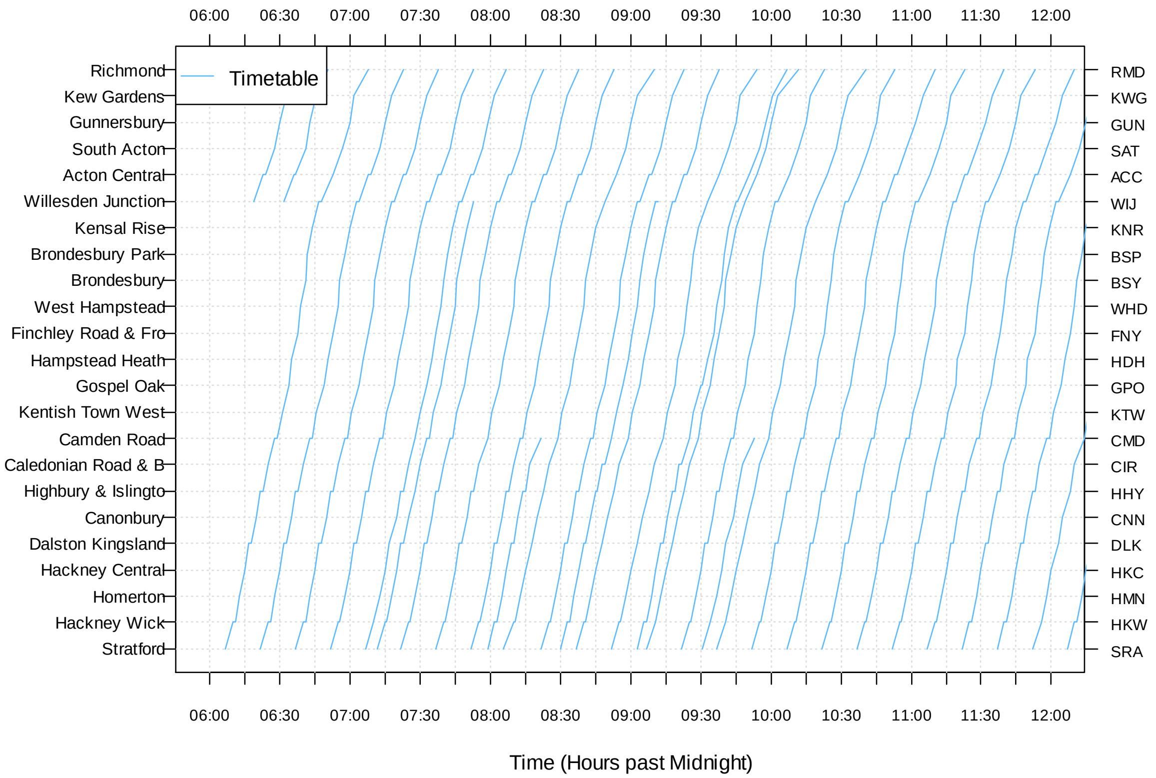

Figure 18: North London Line Westbound weekday morning timetable

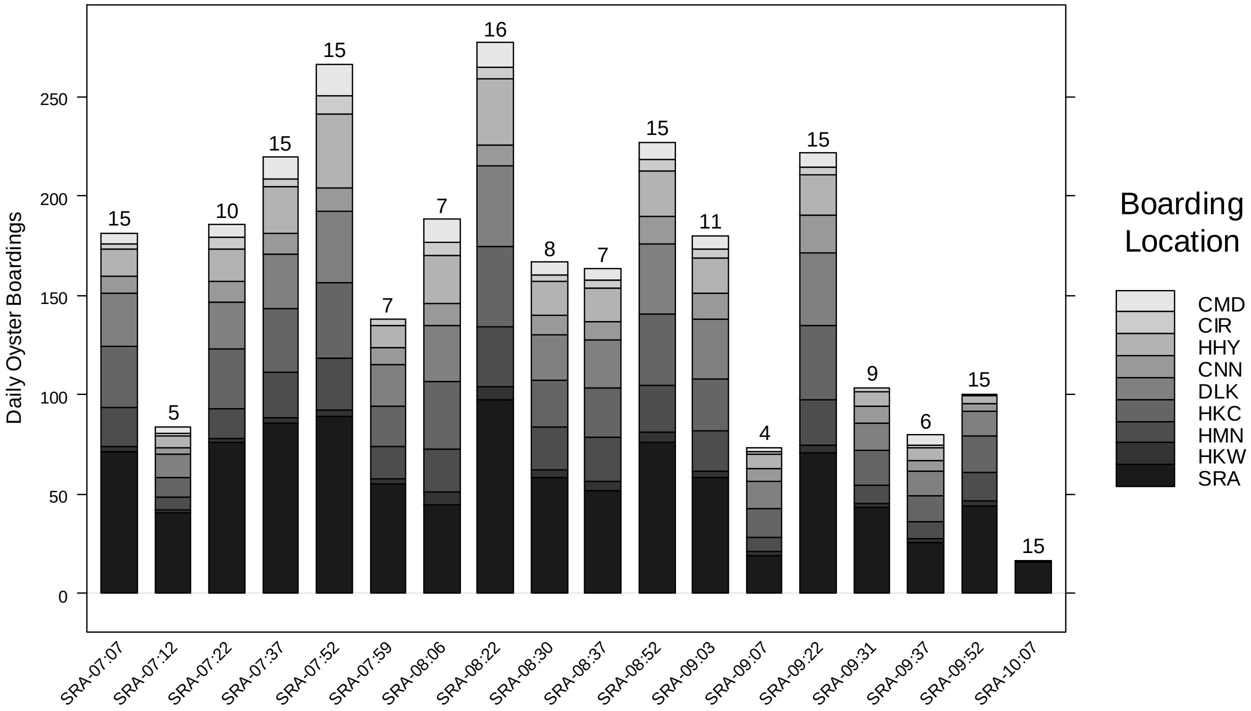

Figure 19: Daily boardings between Stratford and Camden Road on Westbound North London Line services in the AM peak (under hypothetical right-time operations). Each service is labelled atop the respective stacked bar with its leading headway (in minutes).