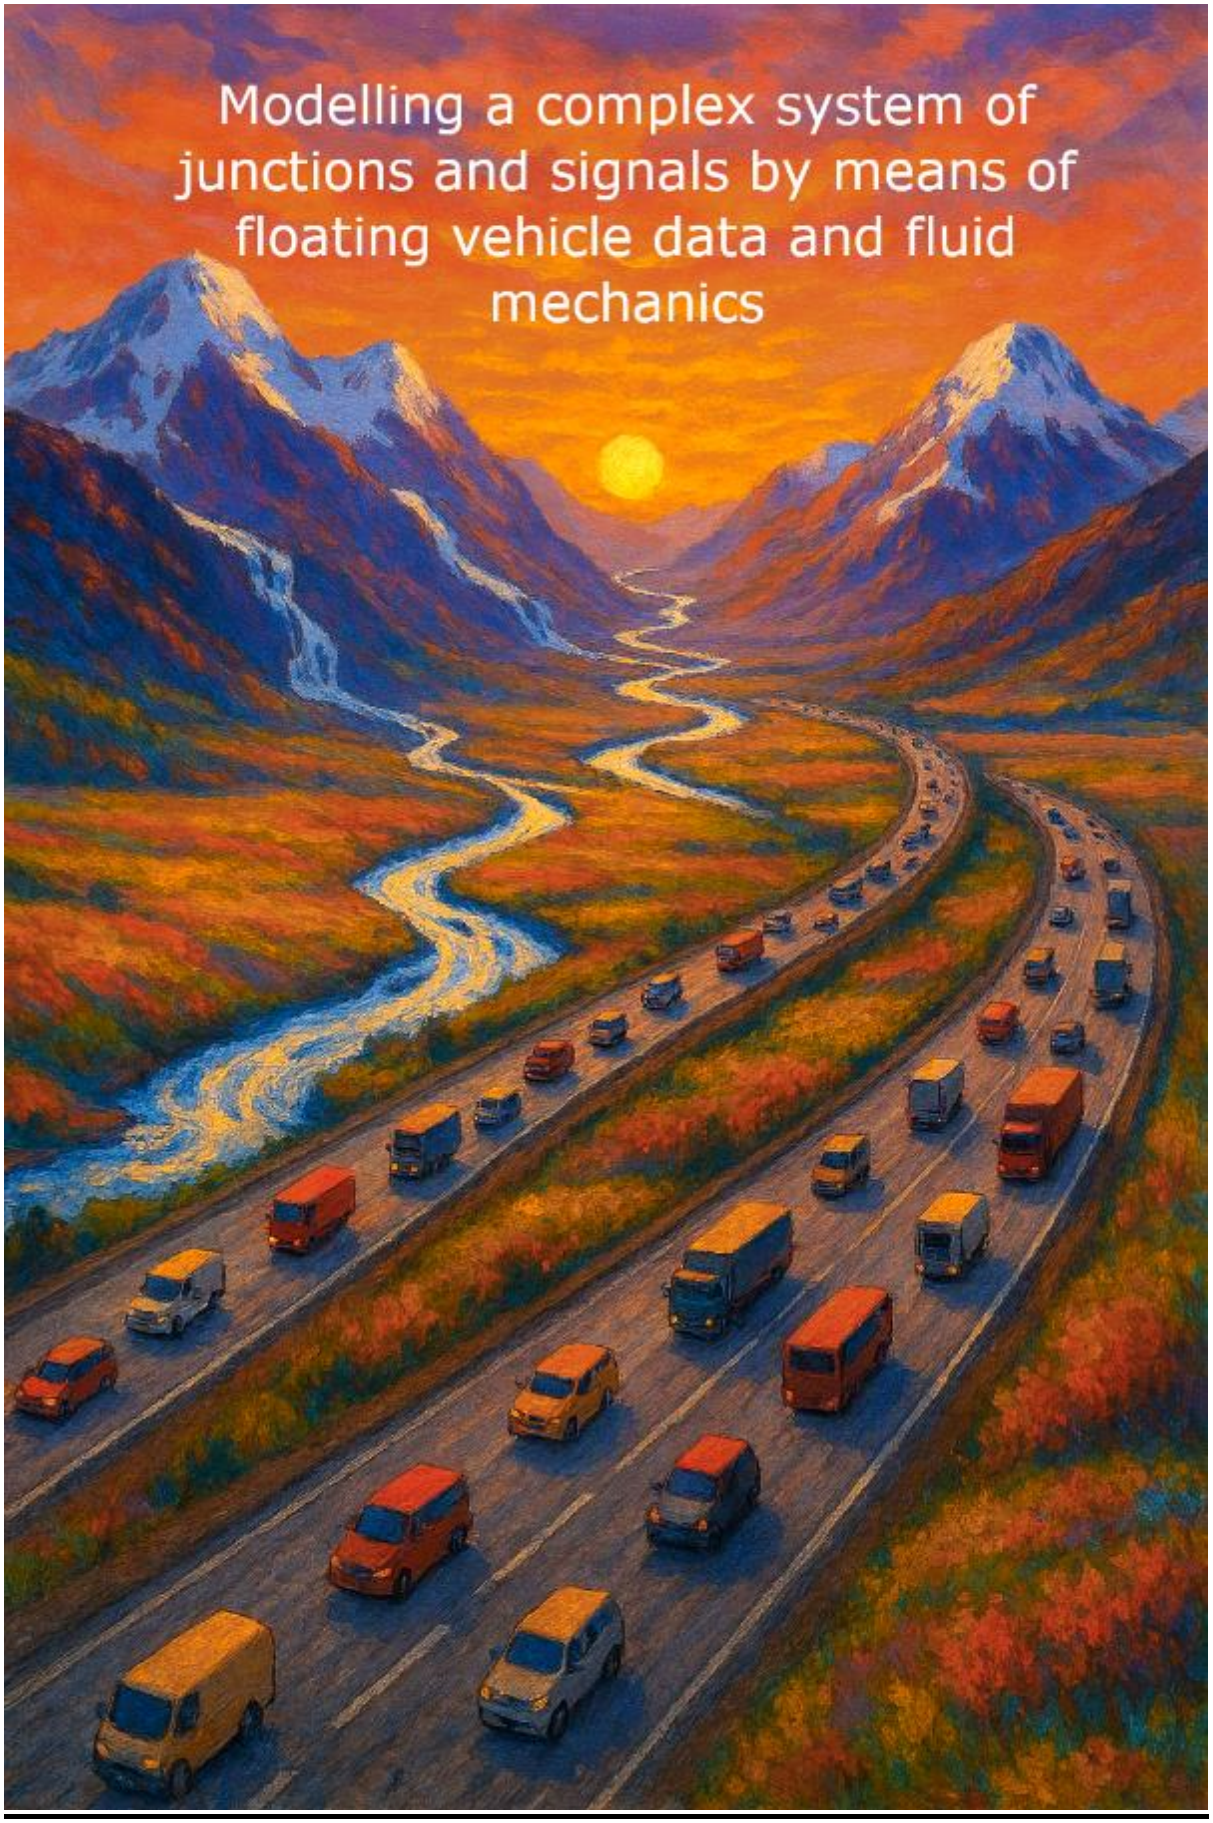

Modelling a complex system of junctions and signals by means of floating vehicle data and fluid mechanics

Modelling a complex system of junctions and signals by means of floating vehicle data and fluid mechanics

Jonathan Flynn	National Highways
Marinos Tsiplakis	Compass IOT Australia
Peter Richards	National Highways

1 Summary
1.1 Theory
1.2 Use Case 1 – Stanstead Airport
1.3 Use Case 2 – Melbourne Airport

2 Theoretical Background

3 Introduction to Floating Vehicle Data

4 Use Case 1 – Stanstead Airport
4.1 Problem Statement
4.2 Getting the data
4.3 Doing the calculations
4.4 Results

5 Use Case 2 – Melbourne Airport

6 Further Investigation Methods

7 Into the future

8 Conclusion

9 Appendices and further information

1 Summary

1.1 Theory

Traffic flow can be modelled at a macroscopic level using the same equations as are used to model the flow of fluids, compressible or incompressible. The mathematical model (dating from 1955) introduced by Lighthill-Whitham-Richards is formulated as a nonlinear partial differential equation derived by the conservation of vehicles in a single road and referred to as the LWR model. This has been developed by other authors over the years.

To collect data in sufficient quantity to feed the models, 'Floating Vehicle Data' will be used. This takes anonymous data from the mobile radio devices (which most vehicles contain, including eCall). It can be processed in quantity by appropriate software.

To demonstrate the scope of the techniques, the use case of Stanstead Airport approach will be considered in depth.

1.2 Use Case 1 – Stanstead Airport

Junction 8 of the M11 is of strategic importance due to the busy traffic flow and nearby Stanstead airport. It has been modified a few times in recent years, but problems still remain. This paper will take a systems approach to the problem, and use floating vehicle data to make the measurements necessary to initiate improvement possibilities. It will be shown that the amount of data remotely collected far exceeds that which could be obtained by direct observation.

By analysing the occupancy and speeds of the different possible routes through the system at different times of day, a picture will be built up.

A quick summary of the floating vehicle data, after eliminating the irrelevant, can be translated into the picture (Figure 1) below.

Key points from the data are that: -

Journeys using the routes shown take 2 minutes longer in peak hours.

Other routes are not massively affected.

The real key to the problem comes from the very bottom of the picture, "Is the rest of the river clear?" Well, evidence shows that it is not! We cannot expect the Jcn8 roundabout system to run well if the traffic is unable to leave the area, and get on its way! It would be wonderful if all the water could simply go over a big cliff, like Victoria Falls!

So, we can repeat the process around Bishops Stortford…!

Floating vehicle data can also be used to look at the other problem areas to a greater level of abstraction. This can include near miss analysis via heavy braking, video analysis, etc.

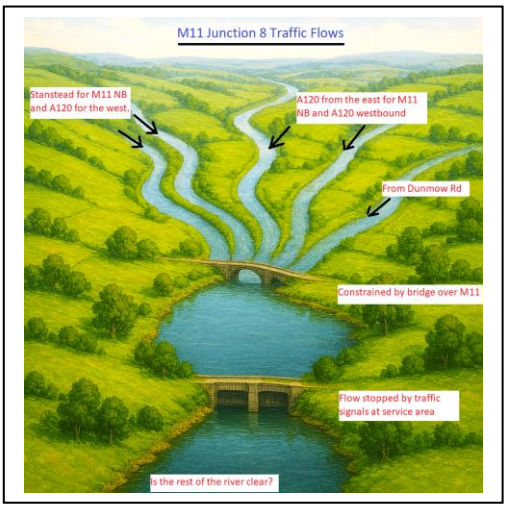


Figure 1

Illustrative diagram of M11 junction 8 traffic flows.

1.3 Use Case 2 – Melbourne Airport

To validate the process, a second use case was chosen, Melbourne Airport in Australia.
The numbers from the Compass system were plugged into exactly the same Excel workbook, with no other changes. This showed: -

a) There was no great traffic problem anywhere! (Result validated from local knowledge)

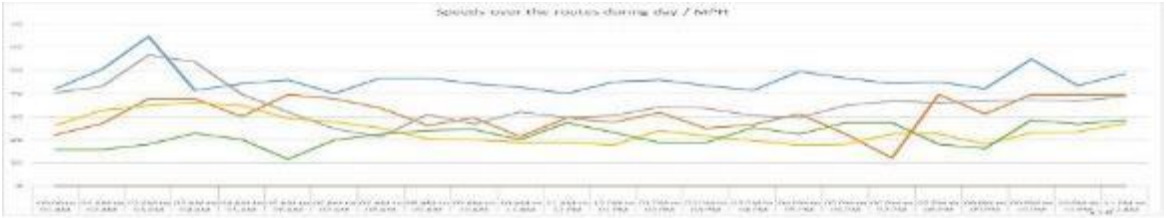
b) The fairly predictable conclusion that the approach roads became busier (and slower) during the working day. Traffic levels could be correlated to flight arrivals and departures.

c) The road **from** the freight centre, Apac Drive, is busiest around 06:00hrs, and the road **to** the freight centre is busiest around 19:00hrs. Journeys in these peak times taking a couple of minutes longer. This must be understood in the context that freight is flown in and out at night, as daytime hours are occupied by passenger traffic.

d) Apac drive has two roundabouts and traffic signals, so it is no surprise that this is the slowest of the journeys investigated.

This use case shows the power of the technique, in that the traffic flow around Melbourne Airport can be investigated from the comfort of an armchair in England!

By the same method, traffic flows can be readily investigated anywhere without leaving the office!



2 Theoretical Background

"We will examine traffic flow in complex junctions in a different way"

The macroscopic model was first introduced by Lighthill and Whitham in 1955 and independently by Richards in 1956. The mathematical model introduced by Lighthill-Whitham-Richards is formulated as a nonlinear partial differential equation derived by the conservation of vehicles in a single road and referred to as LWR model.

$$\partial\rho / \partial t + \partial Q / \partial x = \partial\rho / \partial t + \partial(\rho V) / \partial x = 0$$

Note that x is distance from start, ρ is density, Q is flow and V is velocity,

Meaning:- If density increases over time, then flow must decrease to balance.

The model can also handle 'shock waves', eg red traffic signals!

A similar relationship comes from Bernoulli's Principle, via conservation of energy.

Pressure ≡. Potential energy & Flow ≡. kinetic energy. The sum of both remains constant. So, if **speed increases then pressure drops** (and vice versa).

Flow (Kinetic Energy) is the actual flow rate observed.

Pressure (Potential Energy) is the desire to 'get on with' the journey a bit quicker.

Simple thought experiment… Flow rate is not so good, so drivers get frustrated and apply pressure to continue.. (Sounding horn, tailgating, gesticulating, etc!)

Conversely, when flow rate is good, the 'desire to get on with it' is reduced. It must be remembered that the culture of the people will influence this!!

Pascal states that pressure is applied equally in all directions, (including direction of travel). Charles gas law relates pressure and volume to temperature.

Flow slower implies pressure increases (exaggerated by air temperature!)

Kinetic energy	+	potential energy	=	Constant
Dynamic pressure	+	Static pressure	=	Constant

Highest speeds occur when pressure is lowest.

Temperature may be thought of in the way we look at the culture of the drivers in the area. Calm and collected, or impatient and gesticulating - 'hot headed'. (Varies by region). So, imagine a country full of hot headed impatient drivers. 'Temperature' is higher, so 'Pressure' is higher! (And progress is generally slower!)

We can also think that air temperature getting higher might make drivers more impatient, but with air conditioned vehicles, it is difficult to make the connection.

Bernoulli law states that slower movement exerts higher pressure.

In the macroscopic model, traffic flow is described by comparing the movement of vehicles with the motion of a (compressible) fluid. Here, the traffic flow is represented by density and average speed of vehicles. It may be imagined as a 'helicopter view'.

In the microscopic model, traffic flow is described by the behaviour of each single vehicle in the road. The movement of each vehicle is represented by ordinary differential equation where the dynamic equation depends on velocity and position of the next vehicle. This model is complicated to solve, and leads to discussion of details like loops, stop lines, intergreen times, pedestrian push buttons, etc. Such discussion is outside the scope of this paper.

The above discussion brings many mathematical tools into the toolbox.

The most relevant is the Flow fundamental: Flow = Density X Velocity

Floating Vehicle Data gives velocity. How to get Flow or Density?

a) Use knowledge that 95% vehicles have a mobile device. (sample a loop count vs Floating Vehicle count and get ratio)

b) Use **Vehicles / hour = 3600 / { L/V + 0.7626 α V + 0.681804 α }**
As demonstrated in JCT conference 2024. Same idea as Greenshields model.

c) Any other ideas?

3 Introduction to Floating Vehicle Data

The theory in Section 2 needs measurements to feed it, that is, speeds and journey times across many routes, at many times of day, in enough volume to be statistically meaningful. Traditionally this meant sending people out to count and survey — slow, expensive, dangerous at the roadside, and limited to the days someone happened to be standing there. Floating Vehicle Data (FVD) removes that constraint. Instead of measuring traffic from a fixed point, each vehicle becomes its own moving sensor — it "floats" with the stream and reports where it was and how fast it was going. Aggregate enough of them and the speed and time profile of any route, on any day, can be reconstructed without leaving the office.

Not all FVD is equal, and the difference matters for what follows. The coarser sources — positions inferred from mobile phones on the cellular network, or from apps running on handsets, locate a device to within a rough area, not a vehicle to within a lane, and cannot reliably tell a car from a bus from a pedestrian on the pavement. They are useful for broad-brush movement, but blunt at a single junction. The high-fidelity sources are different in kind, they are traces of actual vehicles, recording where each one truly went and how fast, point by point along its path. It is this latter kind — vehicle trajectory data — that this paper relies on, and it is worth being clear about why.

The contrast that matters is with the aggregated "link speed" products highway authorities will know well. Those condense the journeys of many vehicles into a single representative speed for each road link. It is a convenient summary, but it is a model output, not an observation and the averaging smooths away the brief stops and queues that occur at exactly the points we care about. The result is that aggregated link speeds report a network as flowing more freely than the driver sitting in the queue would recognise. That is fine for estimating a sat-nav arrival time but is of little use for finding the junction that is actually failing, because the failure has been averaged out of the data before it ever reaches you. Raw trajectories make the opposite trade, nothing is smoothed, so every stop and every metre of queue is still in the record. In the fluid analogy, an aggregated link speed tells you the river's average pace from source to sea; the trajectory shows you the eddy forming behind each rock.

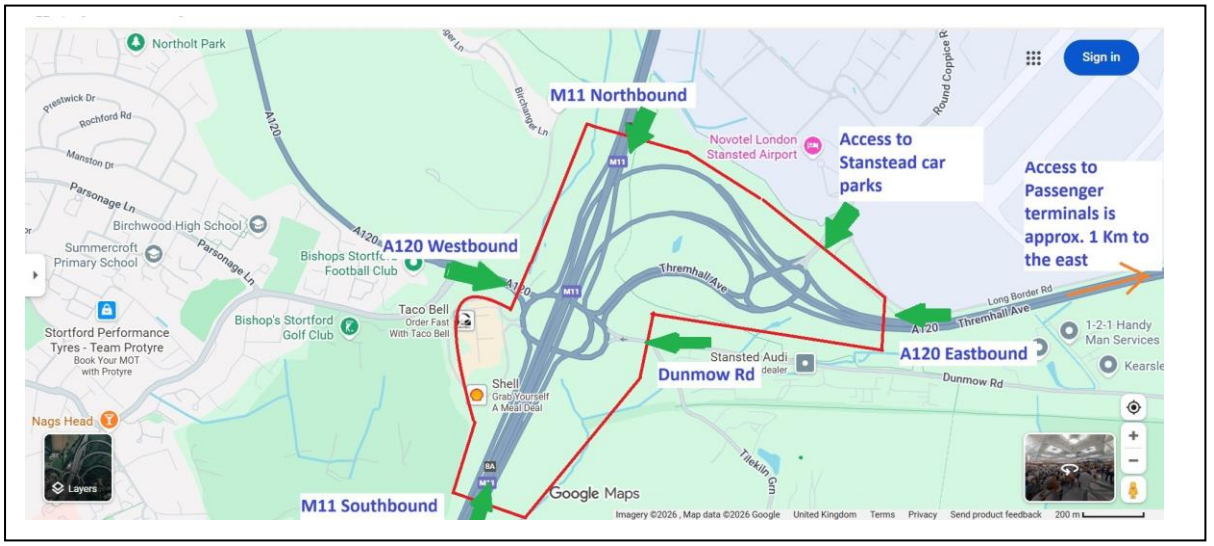
That fidelity buys two things the rest of this paper depends on. First, a true per-point speed profile along each route, rather than an average struck between two widely-spaced fixes — which is what the macroscopic model in Section 2 must be fed if its outputs are to mean anything. Second is depth: the same trajectories also record exactly how vehicles behaved — where they slowed, stopped and queued — so once the model has identified where and when a problem occurs, we can return to those precise spots and look more closely, see Section 6.

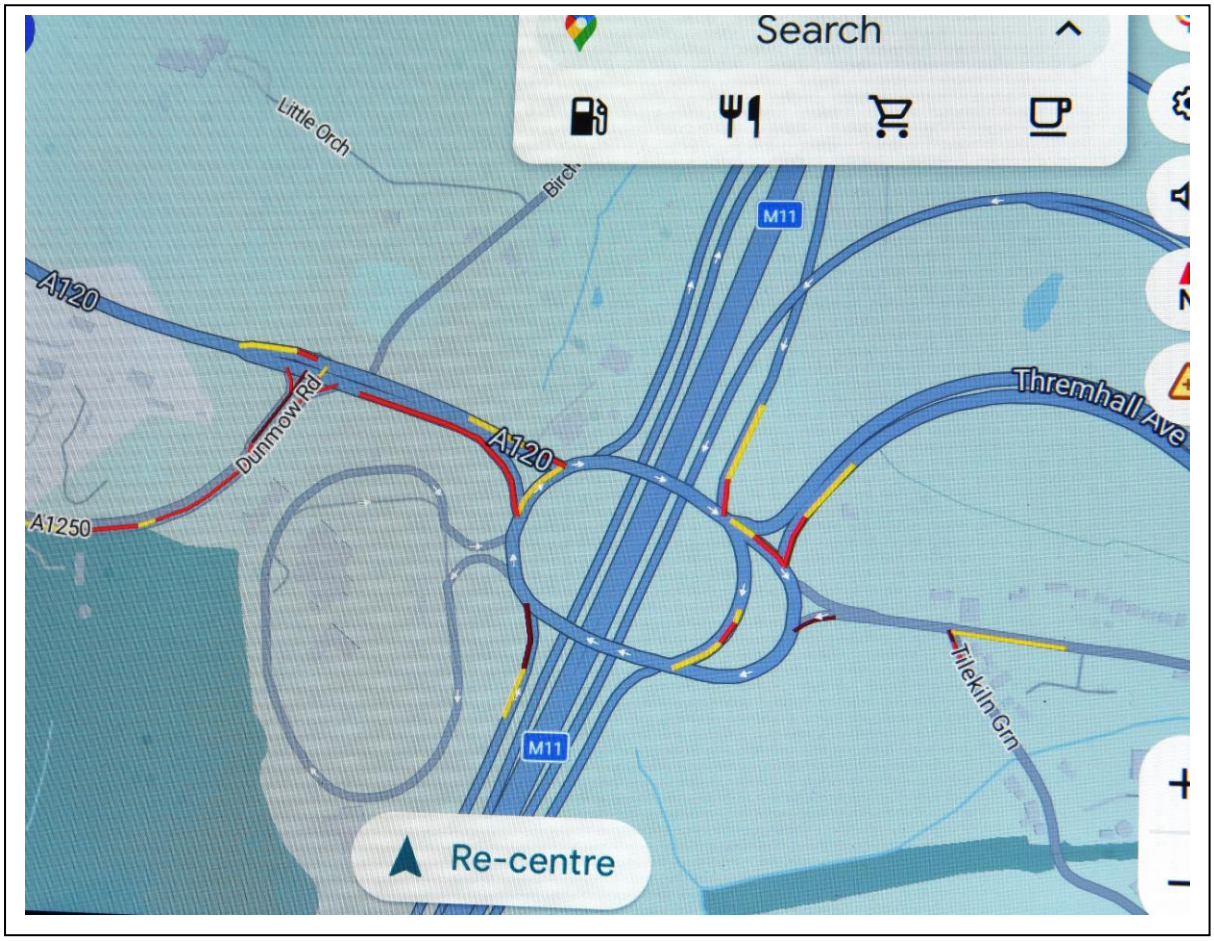
Throughout, the data is anonymous and observational. No individual vehicle or driver is identified; what is analysed is the aggregate behaviour of the stream. We are not tracking one molecule of water down the river — we are measuring the flow of the whole river, which is precisely what the model requires.

4 Use Case 1 – Stanstead Airport

4.1 Problem Statement

The methodology is applicable to any complex system of roads. A ring, rectangle, or other boundary will be drawn around the system. There will be n roads entering the system, and it is assumed that traffic could leave by any of them, (except the one they entered on). For the purpose of this analysis a ‘road’ is only one way. The opposite carriageway counts as another road. (This is reasonable, since it is quite possible that a large circulatory may be used for traffic changing direction after travelling along a dual carriageway). So, with n roads entering the system, there are n(n-1) routes to consider. For each of these, we need to know the trip distance, average time taken, and how this varies over the course of a day. In practice an hourly average will be taken. Hence there will be 24n(n-1) data values to collect and analyse. There will also be another n(n-1) distances to measure. These can all be readily collected from Compass system.


M11 Northbound
Access to Stanstead car parks
Access to Passenger terminals is approx. 1 Km to the east
A120 Westbound
A120 Eastbound
Dunmow Rd
M11 Southbound
Northolt Park
Birchwood High School
Summercroft Primary School
Bishop's Stortford Golf Club
Taco Bell
Shell
Stansted Audi
Nags Head
Novotel London Stansted Airport
Google Maps


Search
Little Orch
M11
A120
Dunmow Rd
A1250
Thremhall Ave
Tilekiln Grn
Re-centre

4.2 Getting the data

It is reasonable to ask where all this data actually comes from — millions of journeys is a great many vehicles, and a model is only as trustworthy as the traffic feeding it. The answer is that no single fleet is the source. The data is aggregated from a broad spread of real vehicles already going about their business on the network, each contributing its anonymous trace:

- Hire and rental fleets — cars from the major rental companies are routinely fitted with telematics and report position and speed as they are driven. These put ordinary passenger cars, driven by ordinary drivers, into the sample.
- Insurer "black box" units — the SIM-based devices fitted under telematics ("pay-how-you-drive") insurance policies. These add a further slice of the private-car fleet, again in normal everyday driving.
- Commercial and logistics fleets — the HGVs and delivery vans run by national hauliers and the large grocery and parcel operators. These are on the road in volume all day and cover the strategic network heavily, ensuring goods vehicles are properly represented and not just cars.
- Connected vehicles — newer cars and goods vehicles reporting natively through their own built-in connectivity, as introduced in Section 3.

The value of drawing from so many sources at once is that none of them dominates. A sample of HGVs alone would tell us about freight but little about the family heading to the airport; a sample of hire cars alone would miss the trucks leaving the service area. Blended, the mix spans the vehicle types, times of day and trip purposes that actually use Junction 8 — which is what lets the journey-time picture in the sections that follow stand as a fair reflection of the whole stream, rather than one corner of it.

With the raw material established and representative, extracting the specific figures Section 4.1 called for is straightforward. Within Compass we draw a boundary around the Junction 8 system, then define each route by its two ends which is simply a point on the entry arm and a point on the exit arm. There is no need to specify the path between. Compass road-matches every real vehicle that made that journey and returns its matched distance and its journey time, split by hour of the day, over whatever historical period and day type we choose. Repeated across the routes, this produces exactly the grid of distances and hourly journey times the calculations in Section 4.3 require without a single cone being put out.

4.3 Doing the calculations

For all routes through the system: -

Distance (length) L is known and is constant for each route.

Time t to cover length L is measured (and is recorded hourly from many vehicles).

Speed, v, over the route is therefore L/t. The set of all speeds on that route is S.

Change of speed over the day is therefore derivative, dv/dt

Average speed over that route during the day is $(1 / 24) \Sigma v$

Since some routes will be short and quick, while others longer and slower, speeds can be normalised thus: (to compare 'like with like')

Normalised speed is $\{ 24 v / \Sigma v \}, \forall v \in S$ which can be plotted in the time domain. This plot should show comparable speeds on all the routes, slowing somewhat during the busy times. However, in the case of problematic routes, there will be even more slowing visible on some traces.

To identify the most problematic cases 'by machine' (as opposed to looking 'to identify trends'), the second derivative of speed is taken, d^2v/dt^2 This is not so simple, as initially the recorded data values do not constitute a differentiable function.

The bigger +ve values correspond to bigger changes in speed, suggesting a bigger problem. The 'machine' can then sort the second derivates by number to reveal the routes with the biggest problems.

By this methodology the process of reducing $\{ (24+1) n (n-1) \}$ data values down to a handful of problematic routes can be automated.

Once the problematic routes are known, drawing them on a map will soon reveal and common points of interest, eg traffic signals, narrow lanes, etc. The times of day will be obvious from the speed / time graphs.

The same method can then be used to a lower level of abstraction to investigate the suspected problems, along with other suitable tools (eg camera images).

It is also worth mentioning that if either of the derivatives are close to zero, that is implying consistent journey times over the day, ie no traffic problems!!

Note:- In algebra the 2nd derivative, d^2v/dt^2 , gives the concavity of a graph.

If $d^2v/dt^2 > 0$, then the curve is upwards. If $d^2v/dt^2 < 0$, then the curve is downwards.

In this case the biggest upward changes are being searched for.

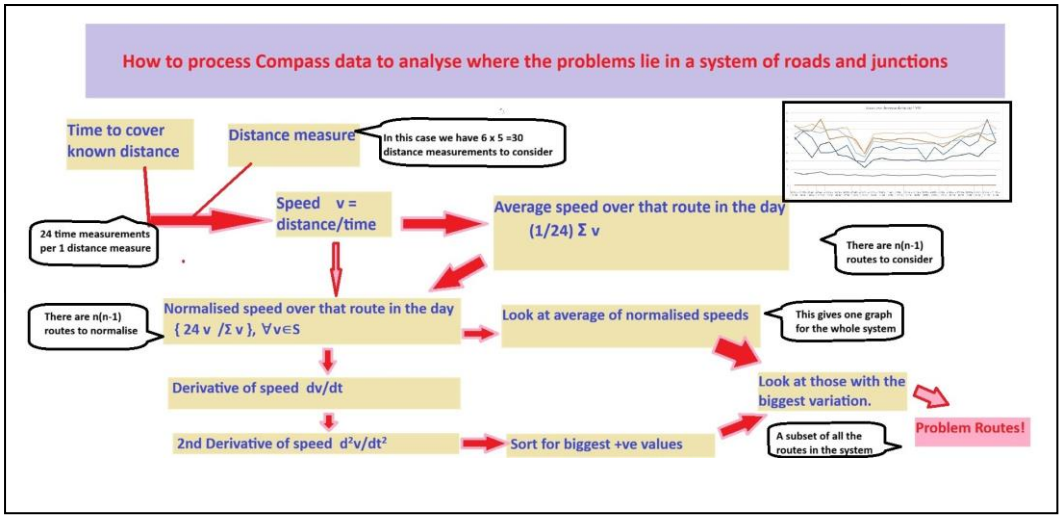


Figure 2 Process Summary

Key Concepts of Differentiability

- **Definition:** A function $f(x)$ is differentiable at a point a if the limit $\lim_{h \to 0} \frac{f(a+h) - f(a)}{h}$ exists.
- **Relationship to Continuity:** If a function is differentiable at a point, it **must** be continuous there. However, a continuous function is not necessarily differentiable (e.g., $f(x) = |x|$ is continuous everywhere but not differentiable at $x = 0$).

4.4 Results – Where to look for the solutions

In the case of M11 junction 8 it can be seen that there are four areas to focus on

a) All the troubled routes pass over the M11 in a westerly direction on the southern most of the bridges. The bridge is three lanes wide, but has to carry traffic from five different feeder routes sorting itself onto a number of outlets, the signing of which may be questioned.

b) All three lanes are stopped for a percentage of the time by traffic signals giving access from the slip road and service area. A waterway which is stopped for half the time has its maximum flow reduced to 50% of its maximum value.

c) The routes westbound into Bishops Stortford must be considered. There is no quick exit (as there is in the Victoria Falls in Zimbabwe). A similar exercise needs to be done in Bishops Stortford to examine the problems there.

d) The traffic flow around the service area needs to be examined in a similar way. Large trucks cannot accelerate quickly, and thought needs to be given to getting them out quickly. It has been suggested that the Birchanger Service Area is used as an unofficial parking area for Stanstead flight arrivals (to avoid airport parking charges). Evidence has been collected to examine that hypothesis, but that will not be discussed here.

In general traffic flow can be considered in the same way as water in a stream. Rocks in the stream impeded the flow, so when the big rocks are removed, the flow improves, but the smaller ones raise their heads… In the same way, the smaller ones get removed, and the flow is even better, but then more very small rocks emerge….

Here the same applies. The big problems direct an examination of smaller ones…..

Some evidence of where to look…..



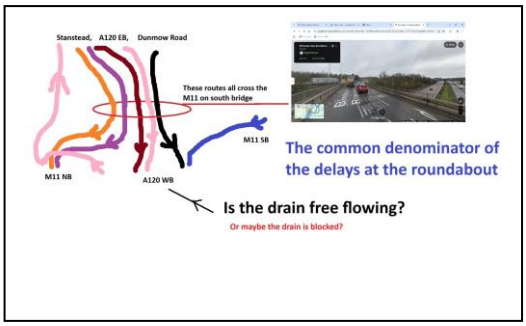
Bishops Stortford at 08:30 AM



Service Area exit around mid day



Service area exit and roundabout



Google Streetview of the bridge

When interventions have been made (eg play with MOVA, move stop lines a few more cms, change intergreen time, etc etc) then repeat the measurement process to see how much improvement there is.

5 Use Case 2 – Melbourne Airport

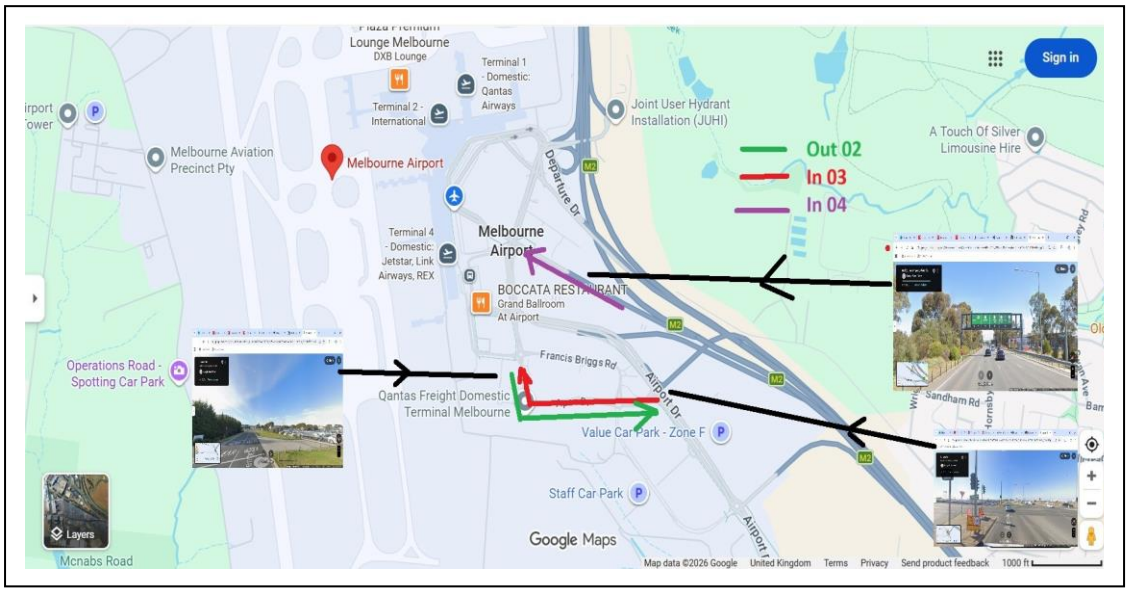
The exact same methodology was applied to Melbourne airport in Australia. Literally different numbers input into the same Excel workbook. The results showed

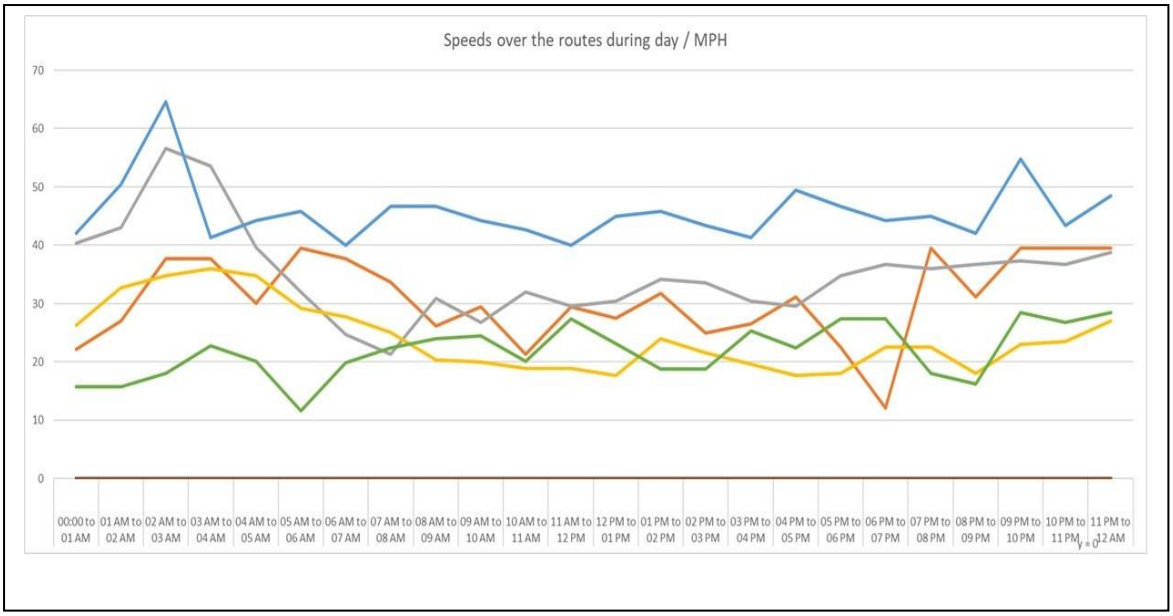
a) There is no 'traffic problem' in the area. Everything flowed reasonably freely at all hours of the day.

b) The road **from** the freight centre, Apac Drive, (Out 02, shown green below) is busiest around 06:00hrs, and the road **to** the freight centre (In 03, shown red below) is busiest around 19:00hrs. Journeys in these peak times taking a couple of minutes longer. This must be understood in the context that freight is flown in and out at night, as daytime hours are occupied by passenger traffic.

c) Apac drive has two roundabouts and traffic signals, so it is no surprise that this is the slowest of the journeys investigated.

As there is no problem to solve, the modelling stopped at this point.





6 Further Investigation Methods

The methods described so far find the problem: they tell us *where* on the network, and at *what time of day*, vehicles are being held up. Once that is known, the same floating vehicle data can be turned to a finer grain to understand *why*, and to test whether an intervention has actually worked. This is the "dig deeper after the first pass" stage, and it is where the richer outputs of the Compass platform come into their own.

A useful illustration comes from the United States, where Dr Zong Tian at the University of Nevada has been using Compass trajectory data as the primary evidence base for signal-timing evaluation (https://www.linkedin.com/feed/update/urn:li:activity:7465114015204052992/). He visualises the trajectories in WaySync, the signal-timing tool he developed, and has used the approach to drive retiming work across corridors in Houston, Orange County and Boca Raton, with the savings on one corridor estimated at around USD$1.7m. The point of interest for us is not the tool itself but the set of measures it surfaces — measures that go beyond average journey time and speak directly to the kinds of problem identified at Junction 8.

Four are particularly relevant.

1. *Time-space diagrams* plot the progression of vehicles along a corridor, making the quality of a "green wave" visible at a glance; this is the natural way to examine coordination through the signalised section of the junction and onward towards Bishops Stortford, where the question is progression rather than delay at any single arm.

2. The *number of stops and the duration of each stop per vehicle* is often a stronger indicator than mean travel time, because it reflects what the road user actually experiences — and it bears directly on the three-lane section that is halted for a percentage of the cycle.
3. *Travel-time reliability*, expressed through the 85th and 95th percentiles and the planning-time index, exposes the tail of the distribution, which is usually where the effect of a signal problem first shows itself; and
4. *Arrivals on green*, or progression quality, the established measure of how well an adaptive or coordinated strategy is performing. Together these allow a problem that has been located at the macroscopic level to be characterised in detail, and importantly to be re-measured after any change to confirm whether it has helped.

A further innovation extends this still further. Compass is now able to pair the connected vehicle trajectories with connected camera feeds, so that the movement of vehicles along those trajectories can be observed visually as well as measured. Where the trajectory data shows *that* vehicles are stopping or queuing at a particular point, the camera view shows *what* is happening to cause it, whether it's the lane discipline on the approach to the bridge, the behaviour at a stop line, or the way large vehicles manoeuvre out of the service area. Bringing the two together turns a numerical finding into something that can be seen and understood and closes the loop between identifying a problem and explaining it.

In every case the principle is the same, first locate the problem with the broad measures, then bring the more detailed tools to bear only where they are needed. The data has already been collected; it is simply a matter of asking more of it.

7 Into the future

The LWR Model as the Engine of the Traffic Digital Twin

The integration of the LWR model finds its most contemporary expression in the concept of a traffic digital twin — a virtual, near real-time replica of a physical transport network. Serving as the core mathematical engine within a digital twin architecture, the fluid-dynamic equations of the LWR framework bridge the gap between raw data ingestion and predictive action.

The Engine of the Digital Twin

In a modern traffic digital twin, high-fidelity data from traffic detection inputs, connected vehicles and roadside infrastructure cameras are continuously streamed into a centralised cloud platform. Using data assimilation techniques built on the LWR model, this substantial influx of point-source data is converted into a continuous, corridor-level view of traffic density and velocity, which can be extended to a network-wide picture when coupled with junction and route-choice models.
Because the LWR model is particularly effective at predicting how bottleneck shockwaves form and propagate backward along a link, the digital twin can run continuous "what-if" scenarios in the background. This allows operators to simulate how a minor incident or sudden lane closure is likely to affect the wider network well ahead of the disruption occurring, giving control room staff time in which to act.

Closing the Loop with Prescriptive Control

This integration ultimately creates a closed-loop system in which the digital twin does not merely monitor the present but actively manages the future. Where the LWR-powered simulation identifies an emerging queue, the digital twin can automatically dispatch optimised control strategies to the physical network. This may include dynamically adjusting signal timings across a corridor, altering variable speed limits, or issuing targeted rerouting advice to connected satellite navigation systems, with the aim of pre-emptively dissolving congestion before it fully develops.
By combining a well-established mid-twentieth-century fluid dynamics model with contemporary digital twin technology, traffic management is evolving from a reactive, descriptive discipline into one that is increasingly autonomous, predictive and adaptive.

8 Conclusion

It has been shown that using fluid mechanics theory, and a macroscopic model of traffic flow, that different insights can be obtained as to complex situations. When it is known that a problem exists, and has been quantified, the same Floating Vehicle Data tools can be used to dig deeper, and investigate further.
Essentially, Compass (or other similar packages) are self service tools to provide insight into this kind of problems. Or reverse the situation, and look at the problem first, and see its effect.
Perhaps most importantly, these investigations can be done from the safety and convenience of the office. No need to go out making surveys, or counting cars!

9 Appendices and further information

Fluid Mechanics principles

Here we consider movement of cars like movement of a liquid.

We need to know

Newton 2nd Law, Force = Mass x Acceleration ($M\ L\ T^{-2}$). Energy Joules = ($M\ L^2\ T^{-2}$).

Pascals Law, Pressure is applied equally in all directions

Boyles Law, Gas Law, pressure inverse relation to volume P1V1 = P2V2

Charles Law, (gas law) Pressure and volume relates to temp. T1/V1 = T2/V2

Bernoulli's Principle, Conservation of energy. Pressure ≡. Potential energy & Flow ≡. kinetic energy. The sum of both remains constant. So, if flow rate is restricted, speed increases and pressure drops (and vice versa).

In detail

Pressure energy + Kinetic energy + Potential energy = Constant along a streamline

Static pressure + Dynamic pressure + Hydrostatic pressure = Constant along a streamline

$P + \frac{1}{2} \rho V2 + \rho g h = \text{Constant along a streamline}$

Where P = Pressure, (Think potential energy)

V = Velocity, (Think speed of flow)

ρ = Density (don't confuse with P)

g = Acceleration due to gravity

h = height (usually same at input & output, thus allowing terms to cancel)

(More complex forms of the equation exist for compressible fluids.)

Flow (Kinetic Energy) is the actual flow rate observed.

Pressure (Potential Energy) is the desire to 'get on with' the journey a bit quicker.

Reference https://www.researchgate.net/publication/227021499_A_Fluid-Dynamic_Traffic_Model_on_Road_Networks

Gabriella Bretti· Roberto Natalini· Benedetto Piccoli

Reference Technische Universitat Dresden Traffic Flow Dynamics 6. The Lighthill-Whitham-Richards (LWR) Model

In microscopic model, the traffic flow is described by the behaviour of each single vehicle in the road. The movement of each vehicle is represented by ordinary differential equation where the dynamical equation depends on velocity and position of the next vehicle.

The macroscopic model was first introduced by Lighthill and Whitham in 1955 and independently by Richards in 1956. The mathematical model introduced by Lighthill-Whitham-Richards is formulated as a nonlinear partial differential equation derived by the conservation of vehicles in a single road and referred to as LWR model.

In macroscopic model, the traffic flow is described by mimicking the movement of vehicles with the motion of fluid. Here, the traffic flow is represented by density and average speed of vehicles

How to use Compass system to extract data:- Here is the link to the quick 2min YouTube video - https://youtu.be/yEsYUsyJnt4

********** **END OF PAPER** **********