PTV GROUP
part of Umovity

CITY OF YORK COUNCIL

Traffic management using a digital twin simulation

Results from a switch-on / switch-off test in York

JCT Symposium, September 2025

1 Transport Digital Twins

Transport digital twins are live, data-driven virtual models of transport systems that replicate the behaviour of real-world networks in real time. They integrate inputs from a variety of sources (such as traffic signals, cameras, radar, GPS and other sensors) to provide a continuously updated picture of how traffic is moving.

There are three definitions within the *digital twin technologies* as summarised in a 2023 report by the Energy Systems Catapult:

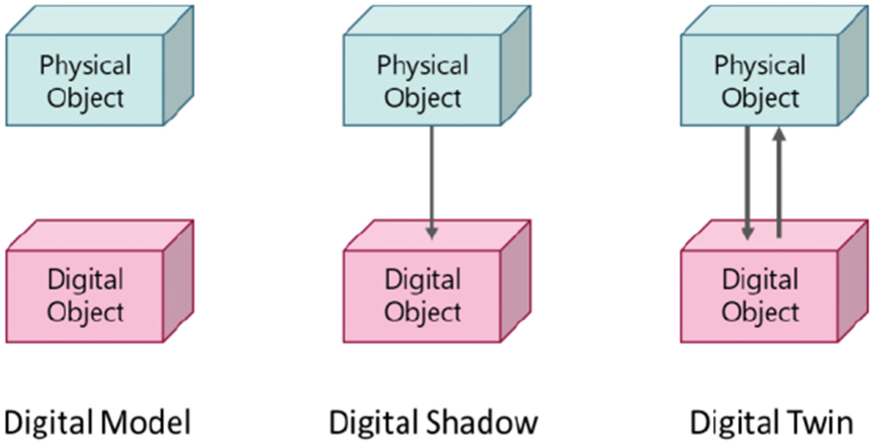

Figure 1: Breakdown of Digital Model, Shadow and Twin

Source: Digital Twins. Model, Shadow, Twin. The Case for Policy Use. Energy Systems Catapult.

We can easily understand a *digital model* as a transport model; a tool that transport planners, engineers and modellers use daily. Whereas *digital models*, or *offline transport models* are calibrated to represent a single previous or future moment in time, a *digital shadow* would be a digital model that is showing the current situation by making use of real-time data. We move into *digital twin* territory, when there is a two-way flow of information between reality and the transport model.

Digital twins are increasingly feasible due to the availability of data on live conditions, such as connected sensors and GPS data, and there is growing interest in their use to address increasing demand on constrained networks at a fraction of the cost of new physical infrastructure. The potential to leverage predictions and simulations makes it possible to automatically test traffic management actions ahead of time and advise network managers on the best course of action.

2 We Need a Robot

In 2021, PTV and the City of York provided a presentation and paper to JCT on the implementation of a city-wide real-time transport model; the first of its kind in the UK. This solution combines a digital model of the city, i.e., a strategic transport model, with real-time data to create a digital shadow, which provides a wealth of information to the city's network managers and planners that they then use to inform their decision making.

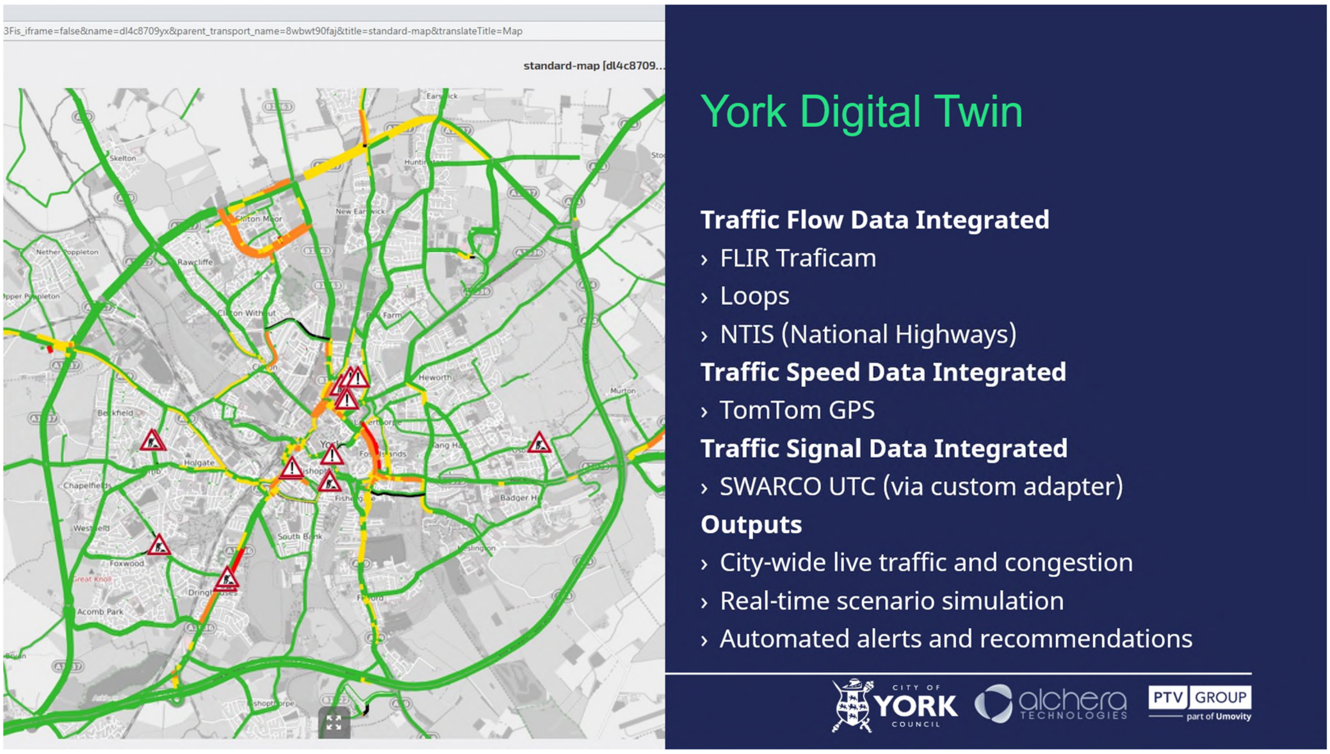

Figure 2: The York Digital Twin

Data regarding traffic conditions provided by the digital shadow have helped to inform short-medium term decision making, such as the planning of works. This has enabled a rapid and regular update to the digital (strategic) model that would otherwise be a costly and time-consuming task.

For real-time traffic management, the city's network monitoring officers use the digital shadow to be aware of situations as they develop, thus allowing them to be more proactive. However, despite having the functional capability within the digital twin, the city has struggled to embed real-time scenario testing and decision support into business as usual. This is due to resourcing, because the system currently requires a human-in-the-loop, and human resources are heavily constrained within local government.

Currently, the human uses the information from the digital shadow to see when and where issues are occurring, before investigating the area on CCTV and using their expert local knowledge to make an intervention in the UTC to change signal plans. They then monitor the situation and make further adjustments as necessary. It is nice to have the additional monitoring capability that the shadow provides, but the real power of the system is in its predictive and simulation abilities. Currently, the human would need to use the GUI to tell the model which plans to test, wait a minute for the results, interpret the results, and then decide on an action to make.

If this process could be automated, then there would not be any need for that additional human resource. The city needed a robot.

3 Making a Robot

For the York digital shadow, a custom appliance had been implemented by Jon Wade at Interoperate Limited, to obtain the timings for every signal on communications in real time. As each cycle completes, the actual signal timings are pushed to the digital shadow to be used within the real-time traffic prediction, which is for the current situation and up to one hour ahead.

We needed a robot that was able to:

1. Monitor traffic conditions
2. Decide when conditions may require alternative signal plans
3. Simulate the impacts of changing signal plans
4. Score the alternative plans
5. Choose the best course of action and implement it

The robot was assembled very quickly, within a couple of weeks, due to the existing city-wide digital twin technologies.

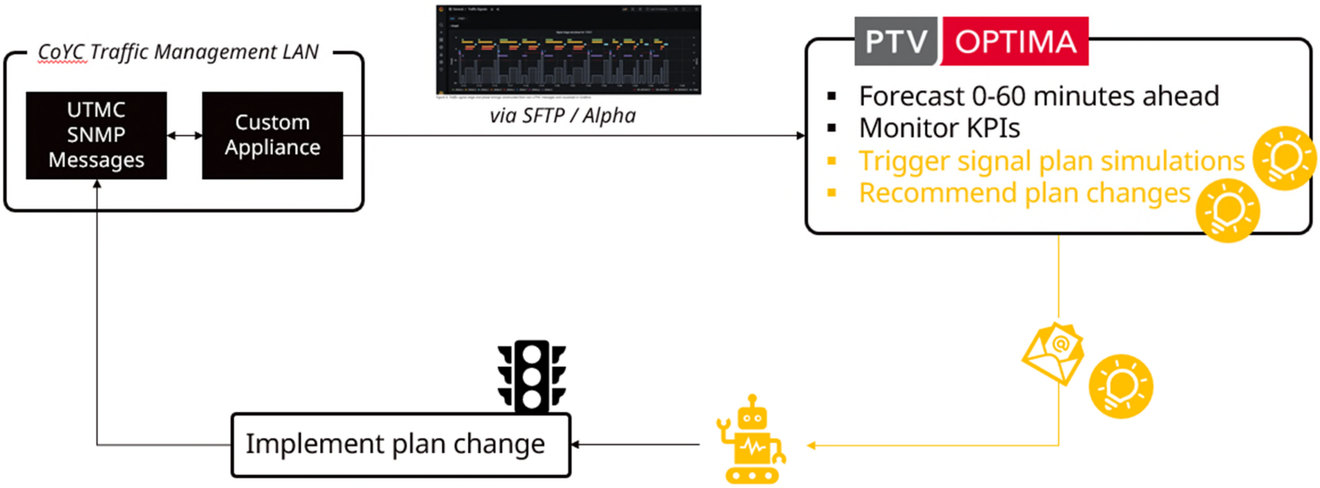

Figure 3: Robot-in-the-Loop

The signalised junction between Clifton Road (A19) and Water End was identified as a real-world test site. This junction sits on a key radial route into the city centre from the north-west, where Clifton Road can become congested at peak times. Water End is the minor road but also has a queue build up during peak times.

With the location chosen, the next step was to draw journey time sections within the digital shadow and begin monitoring those for journey times to determine suitable thresholds for requiring the simulation of alternative plans.

Figure 5 shows the dashboard used for exploring different thresholds, in this case measured as % of travel time delay. Each of the three colours is for one of the three approaches as indicated. 100% means that the journey time is twice that under free-flow conditions, and this was chosen as the threshold for the purpose of this test.

Figure 4: Test Location and KPIs

With a threshold of 100% delay, the bottom half of Figure 5 then shows, for the time-period being explored, how often each of the approaches would trigger simulation of alternative plans. Water End is the main culprit in this case, with Clifton Road southbound also being a trigger, although on a less frequent basis.

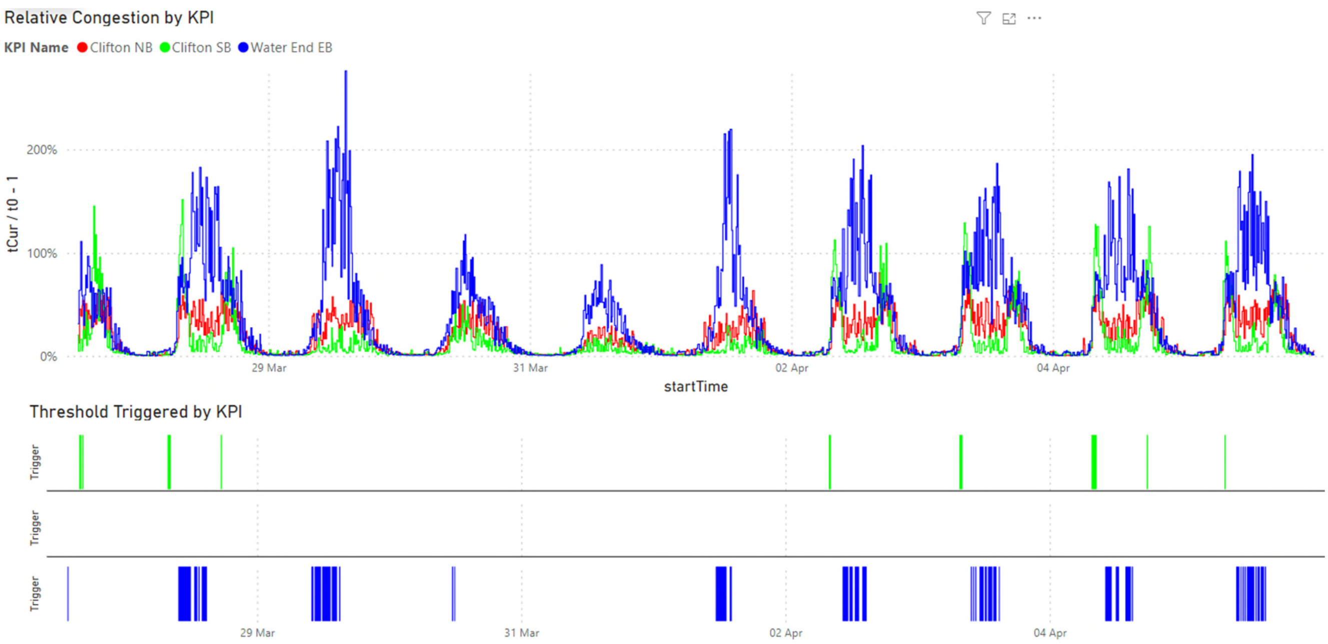

Figure 5: Threshold Design

With the test location and approaches set up for monitoring and threshold alerts, the next step was to create the alternative signal plans inside the digital twin for it use in its simulations. The City of York provided its standard plans for this location in a standard format, which could be very quickly uploaded to the digital twin as shown in Figure 6.

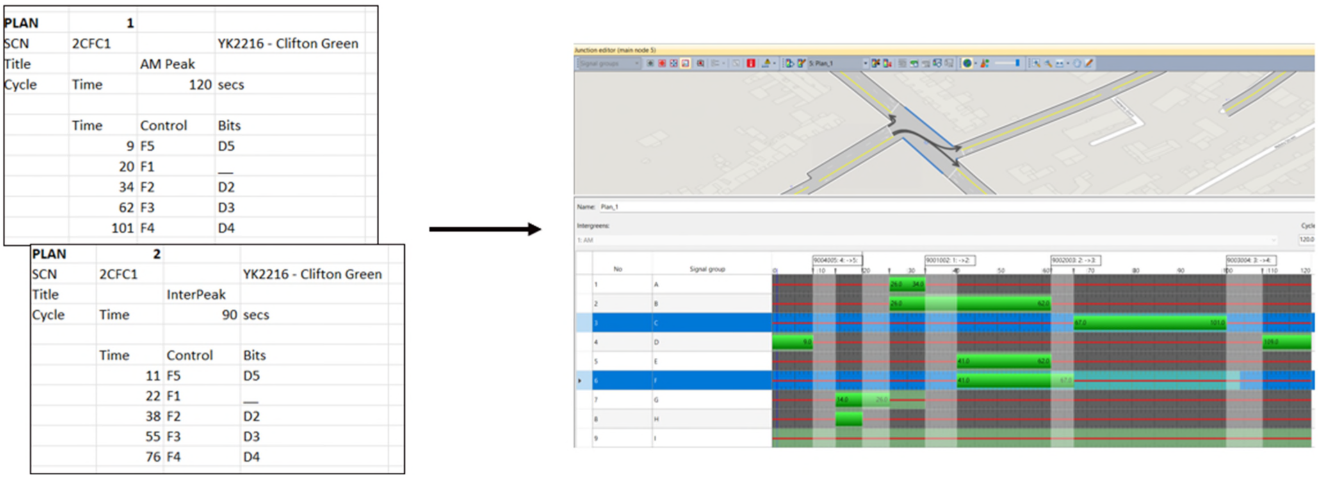

PLAN	1		
SCN	2CFC1		YK2216 - Clifton Green
Title		AM Peak	
Cycle	Time	120	secs
	Time	Control	Bits
	9	F5	D5
	20	F1	__
	34	F2	D2
	62	F3	D3
	101	F4	D4

PLAN	2		
SCN	2CFC1		YK2216 - Clifton Green
Title		InterPeak	
Cycle	Time	90	secs
	Time	Control	Bits
	11	F5	D5
	22	F1	__
	38	F2	D2
	55	F3	D3
	76	F4	D4

Figure 6: Alternative Signal Plans

The alternative plans coded are summarised as follows:

- **Plan 1**: more green time given to Water End
- **Plan 2**: close to the current off-peak plan
- **Plan 3**: the most balanced plan
- **Plan 5**: more green time given to Water End and Clifton southbound
- **Plan 7**: slightly more extreme version of plan 5

Everything was now in place to assemble the robot, which is outlined in Figure 7.

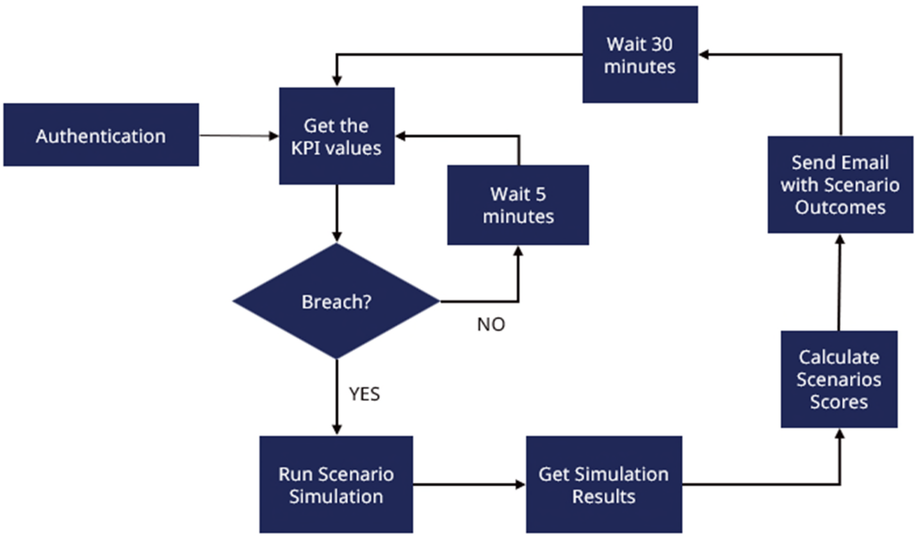

Figure 7: Automation Workflow

This test was set up to check-in every 5 minutes, which aligns with the update frequency of the digital shadow. Every 5 minutes, the robot will check the current journey times on each approach. If there is a breach of threshold, all scenarios will be simulated, which takes about a minute. This simulation takes the current real-time situation, but changes the signal timings for this junction, leaving everything else the same. The results of the simulation are then extracted and compared in terms of the journey time on each approach.

A score is calculated for each scenario, and the signal plan is changed to the one with the greatest score. To avoid lots of jumping around and oscillation, and to allow traffic to settle down after a change, 5-minute automatic monitoring continued, but a 30-minute pause was implemented on allowing additional threshold breaches to trigger a new round of scenario testing.

4 Scoring

The scoring is a fascinating mechanism (honestly), because it has lots of potential for tailoring to policy goals. In the case of this test, the score was kept simple. For each approach, the journey time in the scenario was compared to the journey time of the baseline models for a ‘typical day’. This comparison resulted in a percentage improvement, or worsening of journey time compared to the baseline, for each approach and scenario. The percentage change was then averaged across the three approaches to create a single indicative percentage improvement.

In the case of our test, the aim was to set something up quickly and therefore an even weighting across the approaches was chosen, but it is a trivial change to implement a weighting across the approaches. Furthermore, it would be a simple task to include additional indicators and to weight them into an overall score. Such indicators could include the potential impact on buses, cyclists, pedestrians, and wider network effects.

5 Testing the Robot

The robot was tested by setting up a monitoring system and storing results for two weeks prior to switching the robot on, then two weeks where the robot was active, followed by two further weeks of switching it off.

Figure 8 presents the results of the scores calculated for all five scenarios across two days of the trial. The x-axis is time, and the y-axis is the averaged score in terms of journey time improvement compared to the relevant baseline at that time of day.

Each collection of results at a specific time is the result of the robot deciding that a set of simulations is needed because one of the real-world journey time thresholds had been exceeded. The highest scoring simulation at that time has been circled and the outcome in terms of which plan would be running over the course of the day is shown at the top of the figure.

An interesting finding was that, sometimes, none of the plans were better than the baseline plan at that time of day. In that case, the control room was handed back the reigns to decide on the best course of action.

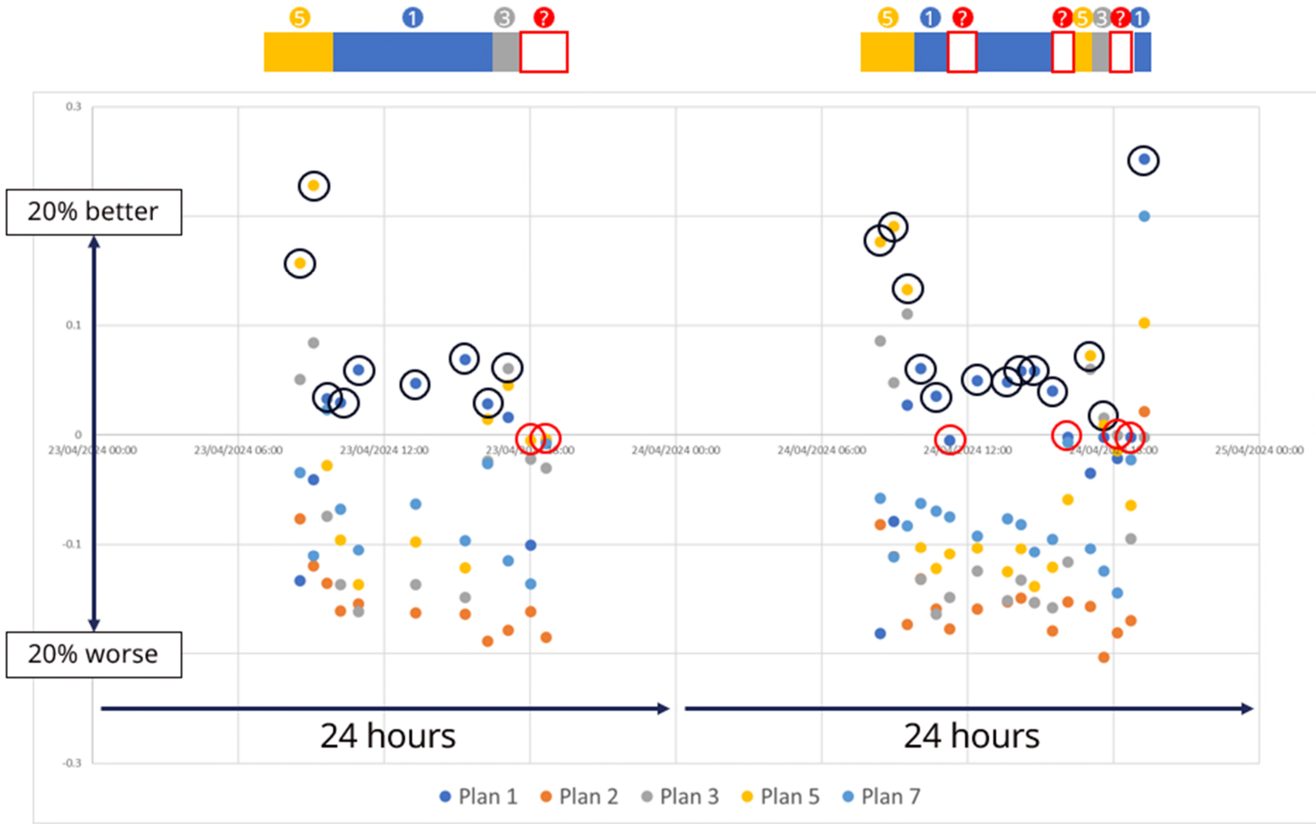

Figure 8: Simulation Scores and Winners

Figure 9 summarises the plans that were selected and implemented via the robot over the course of the trial period. The City of York run their signals either on fixed time, VA, or MOVA. The fixed plan schedule for this junction, which runs VA, is also shown in Figure 9 for comparison.

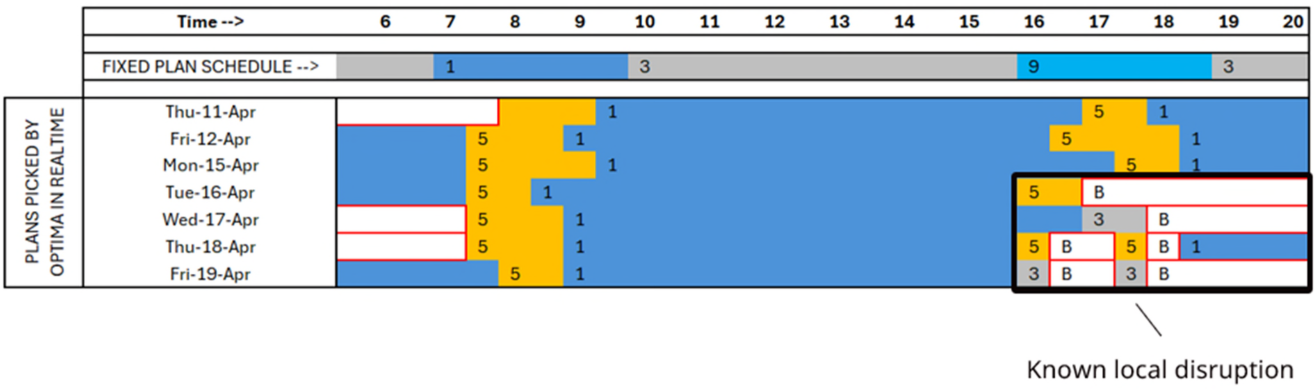

Figure 9: Signal Plans Selected During the Trial Period

The first finding was that the robot did something reasonable; plans changed at peak times. It was interesting to find that during the morning and evening peak, the same plan (5) was chosen. This was the plan giving more green time to Water End and Clifton Southbound. Another interesting finding, which is what you would hope from an adaptive system like this, is that plans changed at different times each day in response to current traffic conditions. The ‘wobble’ at the end of the week in the evening peak was due to local disruption where roadworks started earlier than expected. The reporting to the control room that something unexpected was happening was useful, allowing them to respond proactively.

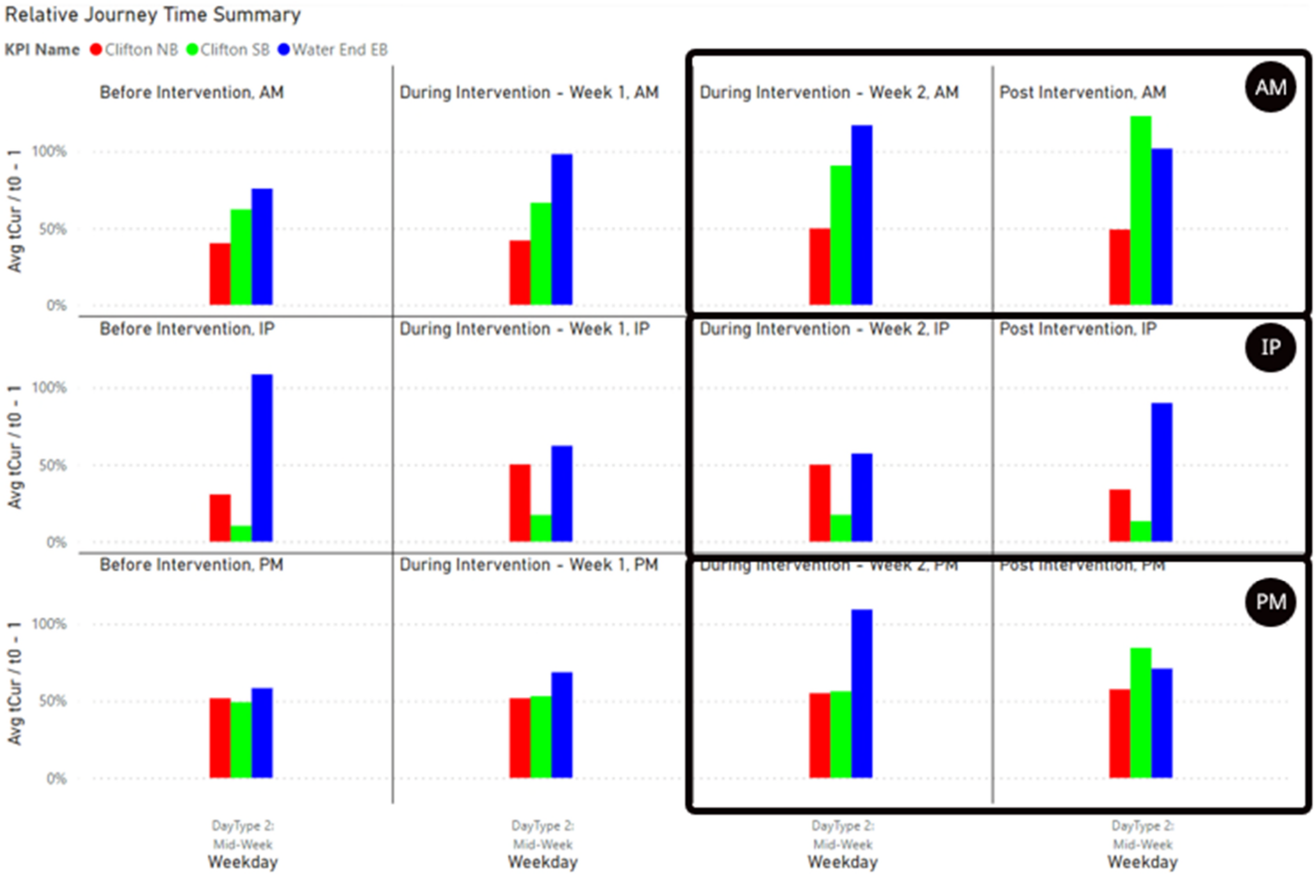

Figure 10: Journey Time Impacts

A benefit of a digital shadow is the large amount of traffic data collected continuously across the city. This allows for impacts to be effectively and easily understood. Due to the test taking place near to school holidays, the most appropriate weeks to compare are highlighted in Figure 10. Each graph shows the average percentage delay compared to free-flow conditions, for each approach (by colour) and time of day (AM/IP/PM). For example, the top-right graph shows the average delay for each of the three approaches in the week following the trial (robot off) in the AM peak.

The results imply:

- **AM and PM**: Clifton Road Southbound was better with the robot, but Water End was slightly worse
- **IP**: Clifton Road Northbound was slightly worse with the robot, but Water End was significantly better

It would be no surprise to a traffic engineer that there would be a trade-off with one approach improving whilst another deteriorates. We were also able to quantify based on the traffic levels the overall impact on vehicle hours which was overall significantly positive (in the order of 30 vehicle hours saved per day).

Equally of interest was to explore journey time variability, as shown in Figure 11. As for the journey time analysis, the two weeks of interest for comparison are highlighted. Each graph shows, for one approach, the maximum, average and minimum 'normalised' journey time over the course of the week. A more 'distributed' plot is suggestive of greater variability.

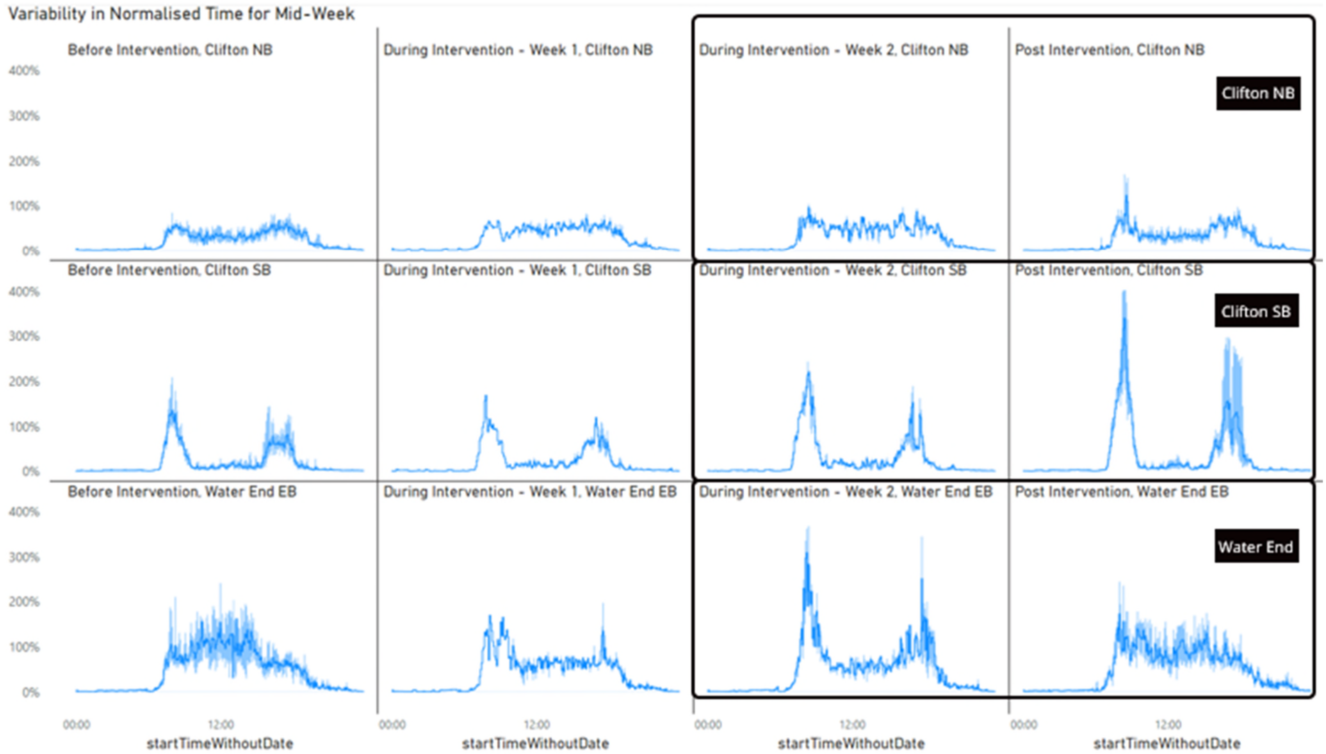

Figure 11: Journey Time Variability

The results imply:

- **Clifton Road Northbound**: relatively little change between the robot-on (left-hand plot) and robot-off (right-hand plot).
- **Clifton Road Southbound**: slightly better journey times with the robot, as expected from the previous analysis, and significantly less variable with the robot in the AM and PM peaks.
- **Water End**: average journey time impacts are as expected from the previous analysis, with potentially also less variability with the robot.

The testing was set up in 2-3 weeks and conducted over a period of 4 weeks only, so one should be careful not to imply too much from the actual impacts. However, the testing clearly demonstrated that the digital twin is responsive, can deliver benefits, and impacts can be well controlled and understood.

The process established here can be scaled and can be tailored to policy goals by designing the KPIs, thresh-olds, and scoring to policy goals.

6 Key Takeaways

York has realised how important the need to leverage data to manage the network is. Data leveraged through the digital twin has provided the ability to quickly identify changes in travel patterns, while also allowing them to understand trends across months and years. It provides the ability to assess strategic management objectives across the wider City network in a timely and efficient manner, given limited resources available. The ability to make a change on one side of the city and see the changes in performance across the rest of network has been fascinating and shown the art of the possible.

Using the digital twin, York has been able to "sweat its assets". For example, using existing traffic detection equipment to produce real-time vehicle count feeds, or pedestrian push buttons to return demand data. Thus, existing infrastructure has contributed to the large-scale data collection, providing York with value for money when installing signal infrastructure.

Optimisation of traffic signal times using real time data was the next logical stage following the success of the modelling process of the digital shadow. This pragmatic solution could be managed and maintained very effectively once the initial concept was identified.

This pilot has shown how quickly innovation solutions can be set up and tested. Moving forward, York supports the use of real time data driven signal optimisation and can see future benefits in continuing the development of this control strategy.

The scoring mechanism provides intriguing potential to not only optimise highway capacity but also to balance it across multiple modes. The need to prioritise active modes evidentially and convincingly is key, as the City looks to show how best to use the highway network.

PTV UK Limited
54 Hagley Road
Birmingham
B16 8PE

+44 (0)121 585 6633
sales.uk@ptvgroup.com

PTV GROUP
part of **Umovity**