

## Lessons from proving ground experiments to investigate junction control

Simon Box<sup>1</sup>, John Lees-Miller<sup>2</sup>, James Snowdon<sup>1</sup>, James Hammond<sup>1</sup>, Andy Hamilton<sup>1</sup>,  
Shashank Gupta<sup>1</sup>, R. Eddie Wilson<sup>2</sup>, Ben Waterson<sup>1</sup>

**Abstract**—An experiment was conducted using the InnovITS proving ground in Nuneaton. Thirty cars with volunteer drivers were asked to drive around a tight closed road circuit causing them to pass repeatedly through a cross-roads junction from all directions. The junction was signalized. In different test-runs of the experiment the traffic lights were controlled by either an automated fixed-time system or by a human using remote control. All vehicles in the test were instrumented using GPS and bluetooth. Video footage from two cameras was also recorded.

The goal of the experiment was to collect data on the performance of human junction controllers. This was motivated by earlier work indicated that human controllers could perform well at this task in a simulated ‘computer game’ environment.

In particular this paper examines some of the issues that arise when trying to simulate an urban road junction in this manner. For example results are presented indicating differences in network performance depending on whether the drivers were instructed to follow a fixed route or a random route of their choice. Thus providing some guidance for maximizing the fidelity of this type of simulation in the future.

The paper also presents a detailed analysis of the sensor data and video footage to measure the performance of the junction under the different modes of control.

### I. INTRODUCTION

#### A. Background

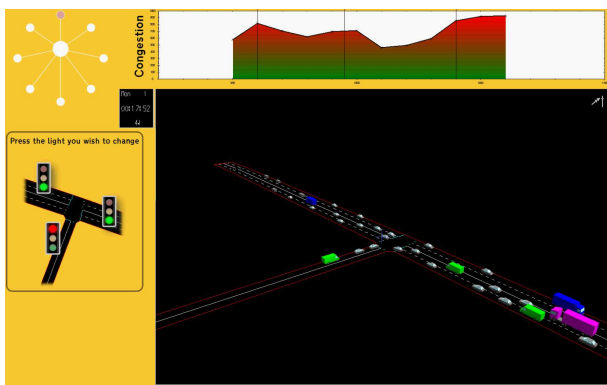


Fig. 1. Screen shot of junction control game.

In previous research humans have been asked to control traffic junctions within microsimulations via a computer

Transportation Research Group, Faculty of Engineering and Environment, University of Southampton

<sup>2</sup>Faculty of Engineering, University of Bristol

game interface, like the one shown in Figure 1 [1]. Video footage of the game being played can be found here: <http://youtu.be/dB3zoGJiODk>. The results of this research have indicated that players of the game *can* be very good junction controllers and even outperform some of the market leading automated junction control systems, such as MOVA [2] and SCOOT [3]. This particular result has prompted the development of machine learning junction control systems that can learn control strategies from human players of the game and from experience [4], [5].

Since obtaining the above results the authors have been keen to collect data on the performance of human controllers in “real-world” tests in a controlled environment. This paper contributes the results of such a performance test. However a controlled environment test is still a simulation of real-world conditions and there can be several sources of infidelity. These are introduced in the following section and examination of these infidelities is the second contribution of this paper.

#### B. Problems of the Experiment

There are two ways that the proposed experiment could be performed. Either a human could take charge of a real road junction, or a junction in the controlled environment of a proving ground.

There are several advantages to using a proving ground for this type of experiment. It is possible to instrument all the vehicles involved, providing rich data with which to evaluate the experiment. It is also possible to control the demand on the junction and standardize this between repeated tests. It is also possible to issue drivers with behavioural instructions.

However a proving ground experiment is a *simulation* of a real world junction and there are some differing characteristics that must be considered. For example, in order to limit the number of drivers and vehicles that are required for the experiment it is desirable to maximise the use of each vehicle by causing them to drive in a closed loop course and pass repeatedly through the junction. To quantify this problem consider that in the experiment discussed in this paper thirty vehicles were used and in each *run* of the experiment, each vehicle passed through the junction on average twelve times. To achieve the same flow of vehicles without a closed loop would require three hundred and sixty vehicles and drivers.

The closed loop course may reduce the fidelity of this simulation, for example it could lead to feedback effects.

When the junction controller releases a vehicle the timing of this decision (weakly) determines when that vehicle will arrive at the junction again. Feedback effects could manifest in a number of ways, including causing vehicles to become ‘synchronised’ with the junction and encouraging ‘platoon’ formation. Such effects are difficult to detect and measure conclusively but efforts were made and these are discussed in Section III.

It is possible that closed loop effects can be managed (or exacerbated) by the instructions that are given to drivers. Specifically, whether drivers are asked to drive a prescribed route or asked to choose their turning movements ‘at random’. Both approaches were investigated in this experiment as described in Section II. Analysis of the difference between these approaches is given in Section III.

## II. METHOD

The experiment took place at the InnovITS “city circuit” test track (<http://www.innovits.com/>), which is part of the MIRA facility (<http://www.mira.co.uk/>) in Nuneaton, UK. Figure 2 shows a map of the city circuit. Cones were placed around the track in order to restrict the vehicles to within the area marked by the boundary shown in Figure 2. Thus a 1-km long figure-of-eight track was formed with a signalised junction at the central node. Cones were also placed where necessary to restrict the width of the road to a single lane in each direction. Thirty volunteer drivers and vehicles were available for the day of the experiment; the mix of vehicles was 25 cars, 2 motorcycles and 3 minibus vans.




Fig. 2. Map of the test track layout.

### A. Track Topology

As discussed in Section I-B, during some of the tests performed we needed to ask drivers to follow a prescribed route around the track. It was considered important to give the drivers clear, easy to follow instructions. So each driver’s route had to be represented by a single line drawn on a map of the track that does not go over any stretch of road more than once **in either direction**.




Fig. 3. Graph of track topology.

It was also important that the set of routes included all possible turning movements through the signalised junction in the centre of the figure of eight. This raises the following question:

“What is the minimum number of routes that include all turning movements while satisfying the constraint that routes must not cover the same stretch of road more than once in either direction?”

Figure 3 shows the topology of the track. On the left there is a graph of the track which has a single node over the junction and four edges that correspond to the clockwise and anti-clockwise routes around each loop in the figure of eight. On the right is an inversion of this graph where the edges in the track graph are now nodes and the node in the track graph is now represented by edges, one for each turning movement through the junction. Our routes can be seen on this turning-movement graph as routes which do not pass through any node labelled **A** or **B** more than once. In this case the routes are represented by the eight simple loops in the graph.

1) *Control Methods:* The signalised junction at the test track is a two stage junction that gives the green light either to vehicles on the North-South running road or the East-West running road. The junction was equipped with an automated fixed time system that switched between these stages every 38 seconds. The junction controller was augmented so that the fixed time system could be replaced by a remote control system enabling human control. The human controller was given a laptop, which was connected wirelessly to the junction controller. Software on the laptop presented the human controller with a very simple control interface of two large buttons allowing the human controller to select the two stages. Also after a stage is selected both buttons are greyed

out for 10 seconds. This was a safety measure to prevent the human controller switching stages too frequently. The human controller was located on a gantry above the junction giving them a good aerial view of the scene. The human controller used in the tests was a layman, that is they were not a professional in the field of transport and had no previous experience controlling junctions either in simulation or the real world.

The experiments described in the following section used both the fixed time controller and human control in different tests. Fixed time control provides a benchmark against which human control performance can be compared it should be noted that most real world junctions use an adaptive control system such as MOVA [2] or SCOOT [3]. Unfortunately these adaptive controllers were not available for this test, although they have been tested against humans in simulation based computer games [4], [5].

### B. Testing Plan

The time available on the day of the experiment allowed for us to give the drivers a short warm-up period where they could drive 4-5 laps of the track and get used to the circuit followed by four 15 minute experimental test-runs with rest breaks in between.

The plan for the four test-runs is outlined in Table I and was as follows: The first three runs were to be under automatic fixed-time control. In test 1 drivers were asked to follow one of the fixed route. These routes were distributed evenly amongst the drivers. In test 2 drivers were asked to pick their turning movements at random. In the test 3 a hybrid strategy was used where 25 of the drivers were given routes and 5 were asked to select their turning movements at random. In test 4 the hybrid strategy from test 3 was used again, but this time the junction was under human control.

TABLE I

SCHEDULE FOR THE FOUR 15 MINUTE TEST-RUNS OF THE EXPERIMENT.

Test #	Instruction	Control method
1	Routed	Automatic
2	Random	Automatic
3	25 Routed + 5 Random	Automatic
4	25 Routed + 5 Random	Human

### C. Sensors

The track and the vehicles were equipped with a number of sensors for monitoring the experiment. Each vehicle was equipped with a **Qstarz Bt-q1000xt** GPS logger. This logged GPS data at 1 Hz and also broadcast a bluetooth MAC address, which allowed it to be detected by six track side bluetooth sensors. There were also two high-definition cameras recording the scene on the track. The results presented in this paper are all derived from the GPS data. Results from the bluetooth sensors are reported separately in [6].

Of the 30 GPS sensors deployed the data from 6 of these sensors has been omitted from the results in Section III due missing or corrupted data. The distribution of routes over the 24 remaining (sensor equipped) vehicles is shown in Table II.

TABLE II

DISTRIBUTION OF ROUTES IN TEST 1, OVER THE SAMPLE OF 24 VEHICLES USED TO GENERATE THE RESULTS IN THIS SECTION.

Route #	1	2	3	4	5	6	7	8
Frequency	3	4	3	4	0	4	3	3

## III. RESULTS

There are number of metrics that can be extracted from the GPS data to evaluate the performance of the junction. Four of these are plotted in Figure 4. These are:

- **Distance Travelled** The closed loop nature of the test track means that the total distance travelled by vehicles during the 15 minute test run period is a proxy for junction performance. The further vehicles travel the less they are delayed during the test. The top-left plot in Figure 4 shows total distance travelled averaged across all vehicles in each test run.
- **Average Speed** Average speed, averaged across all vehicles is proportional to distance travelled, so the top-right plot in Figure 4 gives the same basic information as the top-left plot. This is presented to aid interpretation.
- **Time Stationary** Using the GPS data we can measure the total time that each vehicle spends stationary (defined as below a threshold speed of 1 kph) i.e. waiting at a red light or in a queue. The bottom-left plot in Figure 4 shows total time stationary averaged across all vehicles in each test run.
- **Delay (Gated)** When evaluating the performance of junctions in the real world and in simulations of large networks it is common to consider only a sub-region of the network surrounding the junction. To do this in our tests we implemented a virtual ‘gating’ approach where we consider ‘trips’ of vehicles to begin when they enter a region of 60 m radius around the junction and end when they leave. Travel time for these trips is measured and averaged over all trips. The delay shown in the bottom-right plot in Figure 4 is calculated from the averaged travel time by subtracting the (estimated) free-flow trip time of 15 seconds.

### A. Random vs Routed

With the exception of the delay metric, which will be discussed in Section III-C, all metrics in Figure 4 indicate that the performance of the network is “better” during test 2 than during test 1. However as detailed in Table I, the only difference between these two tests is the instructions given to the drivers. In Test 1 drivers were asked to follow a prescribed route, in Test 2 they were asked to pick their turning movements randomly.

A visual analysis of the video footage from these tests and of animations of vehicle movements generated from the GPS data suggested some possible explanations for this. In test 1 it appeared as if many of vehicles formed into platoons that were making the same turning movement at the junctions, suggesting some synchronization effects. In test 2 it appeared that the opposite was true and that drivers preferred to make

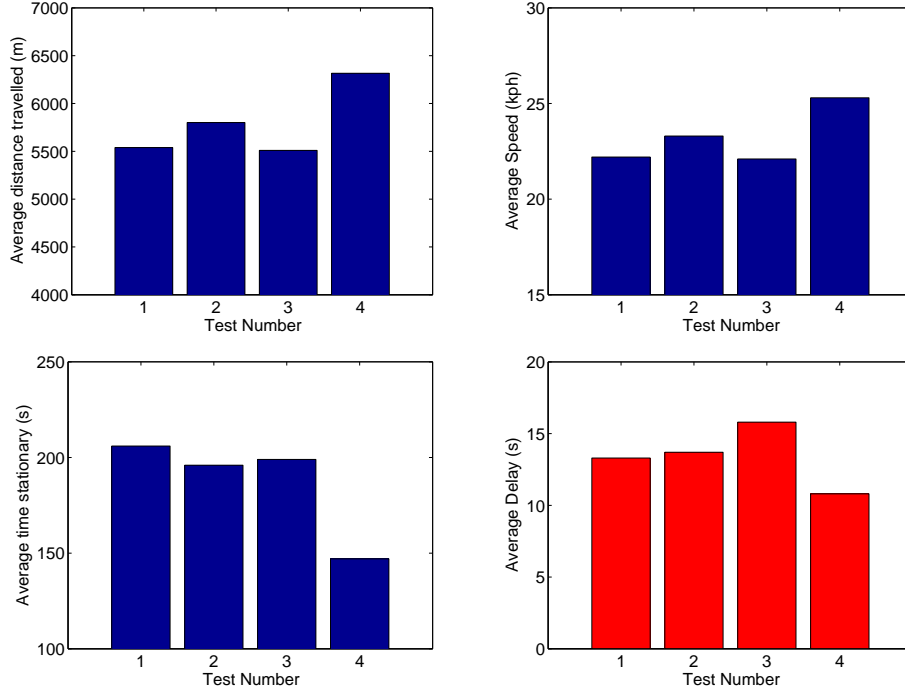


Fig. 4. Four statistics for junction performance and their values in the four tests

a different turning movement to the vehicle in front. It also appeared as if drivers were generally choosing routes to avoid congestion leading to a more homogeneous distribution of vehicles in test 2. It seemed reasonable that drivers in test 2 were (consciously or subconsciously) picking their turning movements in order to minimize their personal experience of congestion and delay. If this was the case we might also expect drivers in test 2 to make less right turns, because of the increased delay associated with this turning movement.

The visual analysis of the data led us to propose three hypotheses to explain the increased performance of the network in test 2. These are:

- **H 1** Vehicles in test 2 were more homogeneously distributed (less platooning) than in test 1.
- **H 2** Vehicles in test 2 were less likely to make the same turning movement through the junction as the vehicle in front.
- **H 3** Vehicles in test 2 made less right turning movements than vehicles in test 1.

We are able to extract evidence from the data to either support or refute these three hypotheses and this evidence is presented and discussed below.

1) *Evidence for Hypothesis 1:* Evidence about how homogeneous or platooned the distribution of vehicles is can be obtained using a *k*-nearest neighbour clustering analysis. At any time-stamp within the GPS data, a snapshot of vehicle positions can be taken. A clustering coefficient can be calculated using (1).

$$R = \frac{\sum_{i=1}^{i=k} (\bar{X}_i) N}{kL} \quad (1)$$

where  $\bar{X}_i$  is the distance from a vehicle to its  $i^{\text{th}}$  nearest neighbour, averaged across all vehicles in the snapshot.  $N$  is the total number of vehicles in the snapshot and  $L$  is the length of the test track.

$R$  will have a value close to 1 if the vehicles are evenly distributed around the track and a value close to 0 if vehicles are tightly clustered into platoons of  $k+1$  (or more) vehicles. The results presented below used a value of  $k = 2$ .

Figure 5 shows the clustering coefficient ( $R$ ) calculated for snapshots at 1 second intervals over test 1 (left-plot) and test 2 (right-plot). The clustering coefficient is low when vehicles are queueing at a red light and when vehicles are driving in close platoons. The plots indicate that vehicles are very clustered at the start of each test. This is because the vehicles were held at an *all-red* signal before the test began, so are sitting in tightly packed queues. As test 1 progresses the vehicles spread out and the clustering coefficient rises to a peak after around 1 minute. The clustering coefficient then reduces as vehicles form queues/platoons and from this point on clustering varies about a stable average. The same thing happens in test 2 except that the stable average is about the same level as the initial peak, suggesting that there is less queue/platoon formation after the vehicles have initially spread out.

The clustering coefficient averaged across all 900 snapshots for each test is shown in Figure 6. This shows that indeed the clustering is higher in test 1 than in test 2. The evidence from these data supports hypothesis 1.

Interestingly the clustering coefficient for test 3 indicates that clustering returns with the driver instructions for test 3, which makes sense because the majority of drivers in this




Fig. 5. k-nearest neighbour clustering coefficient ( $k = 2$ ) for test 1 (left) and test 2 (right). The clustering coefficient is calculated for every 1-second snapshot of vehicle positions and plotted over the duration of the tests.




Fig. 6. k-nearest neighbour clustering coefficient ( $k = 2$ ) for for all four tests and averaged over the 900 1-second snapshots taken for each test.

test are following fixed routes. However platooning reduces again in test 4. Drivers in test 4 have the same instructions as in test 3, but of course the lights are now being controlled by a human, suggesting that the control strategy employed by the human is working to distribute the vehicles on the track more homogeneously. This result correlates strongly with the result in Figure 4, which indicates that the stationary time in test 4 is low, suggesting less queueing.

2) *Evidence for Hypothesis 2:* It is possible to extract the turning movements of vehicles in the tests from the GPS data and hence calculate the probability that a vehicle will make the same turning movement as the vehicle in front of them as they pass through the junction. This is equivalent to the fraction of all turning movements that are the same as the preceding turning movement, which is plotted in the right hand chart in Figure 7. This shows that vehicles were nearly *half* as likely to make the same turning movement as the vehicle in front during test 2. This evidence strongly supports hypothesis 2.

If the turning movements of vehicles were evenly distributed and drawn at random, the *expected* probability of a same subsequent turning movement is  $\mathbb{E}(p) = 0.33$ . This might be a reasonable assumption for Test 2 where drivers are *asked* to select turning movements randomly; and at  $p = 0.358$  the measured probability is close to the expected

value. However in test 1 where vehicles are asked to follow routes we need to take into account the fact that some routes are half the length of the others and their turning movements will be twice as common. We also have to take into account inhomogeneity in the distribution of routes over the sample vehicles (see Table II). Correcting for these two factors, the expected probability of a same subsequent turning movement in Test 1 is  $\mathbb{E}(p) = 0.38$ . While this is skewed slightly high, the measured probability is 21 percentage points higher than the expected value at  $p = 0.591$ .

In Tests 3 and 4 the probability of a same subsequent turns increases as most drivers are now following a prescribed route again. This evidence supports the idea that when following fixed routes the vehicles are forming into platoons of vehicles on similar routes and that this behaviour may be unrealistic.

3) *Evidence for Hypothesis 3:* The number of times vehicles make a right turn can be extracted from the GPS data. The number of right turns, as a fraction of all turning movements is plotted in the left graph in Figure 7, for each of the tests. This indicates that the right turns are approximately equally common in all tests. This evidence indicates that hypothesis 3 is false. Even though drivers in Test 2 appeared to actively select routes to avoid congestion and proximity to other vehicles they did not avoid making right turns.

## B. Human performance

All the metrics for measuring junction performance plotted in Figure 4 indicate a significant increase in performance between Test 3, where the junction was controlled by an optimized fixed time controller and Test 4, where the junction was controlled by a human. Between these two tests average speed is increased by 15%, time spent stationary is reduced by 26% and delay across the junction is reduced by 31%.

It is tricky to extract from the data exactly what it is that the human is doing to achieve the high control performance but we can see from Figures 6 and 7 that an effect of their control is to actively discourage vehicles from forming into platoons of vehicles on similar routes. This is particularly evident in the fraction of same subsequent terms, which is lower under human control.




Fig. 7. Analysis of vehicle's turning movements. The fraction of all turning movements that were right turns and the fraction of turning movements that were the same as the preceding turning movement.

### C. The Delay Metric

The plot for delay in Figure 4 agrees with the other performance metrics in so far as the performance of the human junction controller was better than that of the automated controller used. But on the more subtle differences in performance between Test 1, 2 and 3, the delay metric is different. The main difference between this metric and the others is that it only considers a sub section of the whole network. It might seem that, given that there are no other junctions outside of this subsection, it is reasonable to neglect these areas of the network in the analysis, but these results would indicate otherwise. It appears that the behaviour around the junction is having wider “knock-on” effects throughout the network that are affecting overall performance.

## IV. CONCLUSIONS

The results certainly indicate that, in this case the human controller was able to control the junction significantly ‘better’ than automatic fixed time control for all performance metrics tested. However a full characterization of how well a human can control the traffic would require repeated tests on different junctions, with different controllers and benchmarked against better automated controllers (e.g. SCOOT and MOVA).

Results for Test 1 (where vehicles followed fixed routes) indicate that platooning of vehicles on similar routes was high relative to the other tests. While it is not clear whether this was higher than “normal” and caused by feedback effects, it was the perception of those observing the test that platooning appeared “unnaturally high” and the same subsequent turn probability in Figure 7 is significantly higher than the expected value for a random process.

Results from Test 2 (Where vehicles instructed to make random turning movements) indicate that vehicles were less platooned than when following fixed routes. Also the same subsequent turn probability is close to the expected value for a random process. This suggest that the drivers were quite good at following the instruction to “select turning movements at random”.

Results from Test 3 (Where a hybrid strategy of 25 Routed and 5 random vehicles was used) showed similar levels of platooning as Test 1, although the subsequent turn percentage was slightly lower.

Results from Test 4 (Where a human controlled the junction) indicated good performance but it is not exactly clear why this is the case. Figures 6 and 7 indicate that in test 4 the control of the junction acted to prevent platoon formation and distribute vehicles more homogeneously around the track.

The results in these tests (especially the same subsequent turn probability in Figure 7) suggest that instructing drivers to select their route at random is preferable to asking them to follow a prescribed route. A hybrid strategy is also possible but probably with a higher fraction of random drivers than used here. Other options for routing drivers (e.g. dynamically instructing turning movements by wireless communication) are, of course, possible but were not tested here.

Finally small differences between the delay measured across the junction and other performance metrics that considered the network as a whole may suggest that the practice of excluding areas of the network in performance evaluation, even if they contain no junctions, can skew the results.

## REFERENCES

- [1] S. Box, “What can traffic lights learn from us?” *Traffic Engineering and Control*, vol. 52, no. 8, 2011.
- [2] G. Vincent and J. Peirce, “‘MOVA’: Traffic responsive, self-optimising signal control for isolated intersections,” *TRRL Research Report*, vol. RR170, 1988.
- [3] P. Hunt, R. Bretherton, D. Robertson, and M. Royal, “Scoot on-line traffic signal optimisation technique,” *Traffic Engineering and Control*, vol. 23, pp. 190–192, 1982.
- [4] S. Box and B. Waterson, “An automated signalized junction controller that learns strategies from a human expert,” *Engineering Applications of Artificial Intelligence*, vol. 25(1), pp. 107–118, 2012.
- [5] —, “An automated signalized junction controller that learns strategies by temporal difference reinforcement learning,” *Engineering Applications of Artificial Intelligence*, vol. Online ahead of print, 2012.
- [6] J. Lees-Miller, S. Box, and R. E. Wilson, “Hidden markov models for vehicle tracking with bluetooth.” in *13-3032 TRB Highway Traffic Monitoring Committee (ABJ35)*, 2013.