

Strategic Priority – Buses, Trams and Fire

Ben Hallworth BSc (Hons), MCIHT
ben@brhallworth.co.uk
BR Hallworth Ltd

INTRODUCTION

This paper discusses the benefits which can be achieved by providing priority to selected vehicles, principally in coordinated fixed time networks. The 'strategic' aspect of the paper relates to the fact that the provision of priority in such networks often benefits from a 'plan of action' (strategy) in order to effectively target particular elements of delay.

Fixed time coordinated networks can arguably be among the most 'strategic' of networks, as they can be the result of much thought, development and 'policy' constraints. But whilst they can be perfectly satisfactory for the control of general traffic, they can make it difficult to provide effective priority to vehicles such as buses which do not run at the same cruise time, and/or which do not enter the network in parallel with this general traffic.

In recent years the STM priority system (previously referred to as 'SPRUCE' in West & South Yorkshire) has been developed to allow fixed time networks, and even isolated Vehicle Actuated junctions, to have their timings manipulated for the provision of priority to selected vehicles. Over time the application of STM is becoming more generic for relatively standard situations, but networks with several signalized nodes in them can be anything but 'standard', and can often require priority strategies to be tailored to the particular circumstance.

This paper is split into 3 sections covering examples of the use of STM for 3 classes of priority vehicle, in order of increasing priority needs: **1) Buses, 2) Trams & 3) Fire Appliances.**

STM BACKGROUND

'STM' (standing for Strategic Traffic Management) started as the software-based Priority Tool originally developed by Leeds City Council under the DfT sponsored UTMCO1 project. Its development in both Leeds and Sheffield has been documented over the years in JCT Symposiums¹⁻³, and it is now marketed by *telent*.

STM is currently implemented in 5 authorities: Leeds, Sheffield, Calderdale, Bradford and recently Edinburgh. Priority has now been implemented at over 400 signals, and has been developed to the point that it now has a comprehensive capability including:

- Can monitor all controller reply bits as a basis for making a range of informed timing interventions (or storing for historic data collection),
- Can override normal UTC plans with different synchronized plans (e.g. 'priority' or 'compensation' plans),
- Can alter timings in a variety of different ways ('extensions', 'recalls', 'holds', 'short term offsets', 'cumulative offsets' etc.),
- Can 'pick up' controllers from VA, and provide priority timings via UTC before 'dropping back' to VA,
- Can insert 'advance' controller demands (or similar bits) from detector points upstream of local detection.

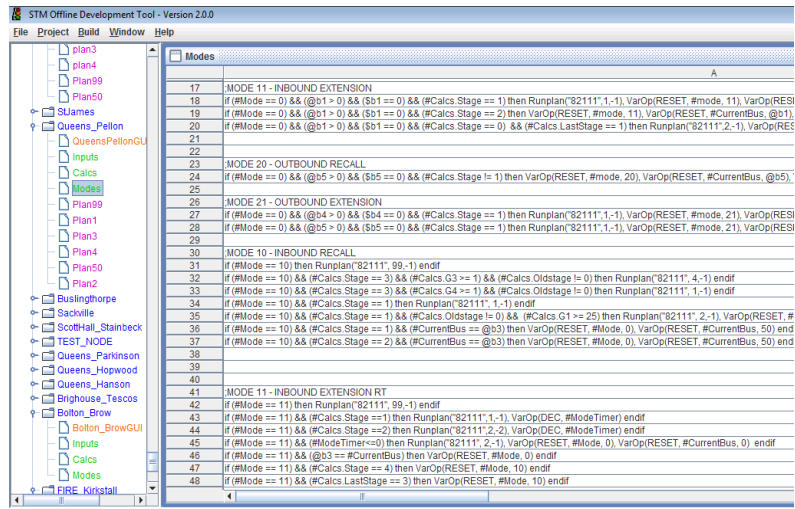
STM ARCHITECTURE

The architecture of the developed version of STM is based around an SQL database. This stores data required to run: node configurations, plans, AVL (Automatic Vehicle Location) detector inputs, and the programmable logic that drives the system. The system is synchronized to the UTC computer, and when controlling signals on-street it monitors detector inputs, processes logic, runs plans, and outputs the necessary stage forces every second. There are now interfaces to both Peeks and Siemens UTC systems.

1. Graphical User Interface (GUI)

A GUI is provided to enable the user to both monitor and influence STM run-time operation. Each GUI is created by the user and can show any parameter available within the STM logic.

Example of STM logic



2. Off line Development Tool (ODT)

The ODT is used entirely off-line to configure plans and related logic sheets. Users are provided with a 'cell based' language to configure plan selection logic, which includes Boolean logic, time-related and plan-related functions. Users can also define their own functions where repetitive logic elements are required.


3. G-bit logger & Journey Time Tool

These are external pieces of software capable of processing UTC reply bits and AVL data from the STM SQL Database. They are key elements in developing the priority strategies.

G-bit logger showing stage green durations at 3 nodes, each row shows a cycle, each green cell indicates an instance of priority

The screenshot shows a G-bit logger interface with three columns for nodes 09412, 09412, and 09413. Each column contains a grid of data representing stage green durations. The data is organized into rows for each cycle, with green cells indicating instances of priority.

Journey Time Tool showing journey times between two AVL detectors throughout the day. Clear banding can be seen in the AM & Peak peaks due to fixed time coordinated network.



1) Buses – West Yorkshire Priority Project

In 2009, West Yorkshire Passenger Transport Executive (Metro) initiated the implementation of a large scale priority project across West Yorkshire - the 'Traffic Light Priority' (TLP) project – with initially 197 junctions (recently increased to 240). Intelligent priority systems were employed to gain the scheme benefits, with AVL virtual detection the input source of choice (virtual detector outputs are here referred to as 'triggers').

Of the 240 sites in the West Yorkshire scheme, sites deriving priority via STM constitute about 80% of the total. These sites require focused thought leading to a range of strategies which are tailored to the specific circumstances of the individual networks. A minority of the sites in this scheme have priority provided by SCOOT, MOVA or local controller priority.

Monitoring of results

As part of the scheme a monitoring report was produced to demonstrate the effectiveness of the traffic signal priority. The report contained 17% of the original 197 total (34 sites). It examined the tested benefits of sample sites, and drew conclusions regarding the results. In particular, it addressed the key project metric that the scheme Benefit/Cost Ratio (BCR) should exceed a value of 2.

The results of the report were based on a 'before and after' protocol developed by the Scheme Working Group, which addressed a number of issues including 'use of triggers for monitoring', 'the sampling rate of junctions', 'the required number of trigger matches' and 'comparator before & after periods'.

NORTHGATE, HALIFAX - STRATEGY EXAMPLE

The bus priority example given here is a group of 3 coordinated junctions in Halifax, operating at a fixed cycle of 80s, and provided with STM priority. This group resulted in an overall bus improvement sufficient to place it around the mid-range of the Benefit/Cost Ratios of the junctions monitored as part of the scheme. However, this was achieved by applying different priority 'tactics' to each junction in the group.

The PM route out of the bus station (northbound from junctions 1 to 3) was seen as a particular delay issue and was specifically targeted in the priority strategy. In the absence of priority, northbound buses frequently fail to coordinate through Junction 3, due to lack of capacity at that junction.

The priority strategy, which provides automatic *compensation* on following cycles, involves:


- **Junction 3.** Providing a *green extension* to improve the northbound bus coordination
- **Junction 2.** Providing a *green recall* complementary to the priority at Junction 3 – this increases the use buses can make of the extension at Junction 3, and also reduces the amount of general traffic running in front of buses.
- **Junction 1.** No action outbound.



For monitoring purposes triggers were located at the entry to Junction 1 and at the exit from Junction 3. By matching pairs of 'entry' and 'exit' triggers for each bus, it was possible over a period of days to build up a meaningful picture of travel times through this network. The chart below shows the resulting scatter plot of travel times for both the 'before' situation (blue dots), and the 'after' situation (red dots).

The scatter plot generated from the Journey Time Tool tends to indicate two horizontal bands where travel times are clustered – the lower band (close to the 50s line) corresponding to buses being coordinated through the first cycle at Junction 3, and the upper band (close to the 100s line) corresponding to buses which miss the first green and have to wait for the next cycle. There are significantly more 'after' (red) triggers in the lower band - this representing a mean improvement in PM travel time of 17s. Were it not for the fact that priority has to be limited to cycles when there is no pedestrian demand at Junction 3 (little more than 50% of cycles), the PM benefit would be significantly higher

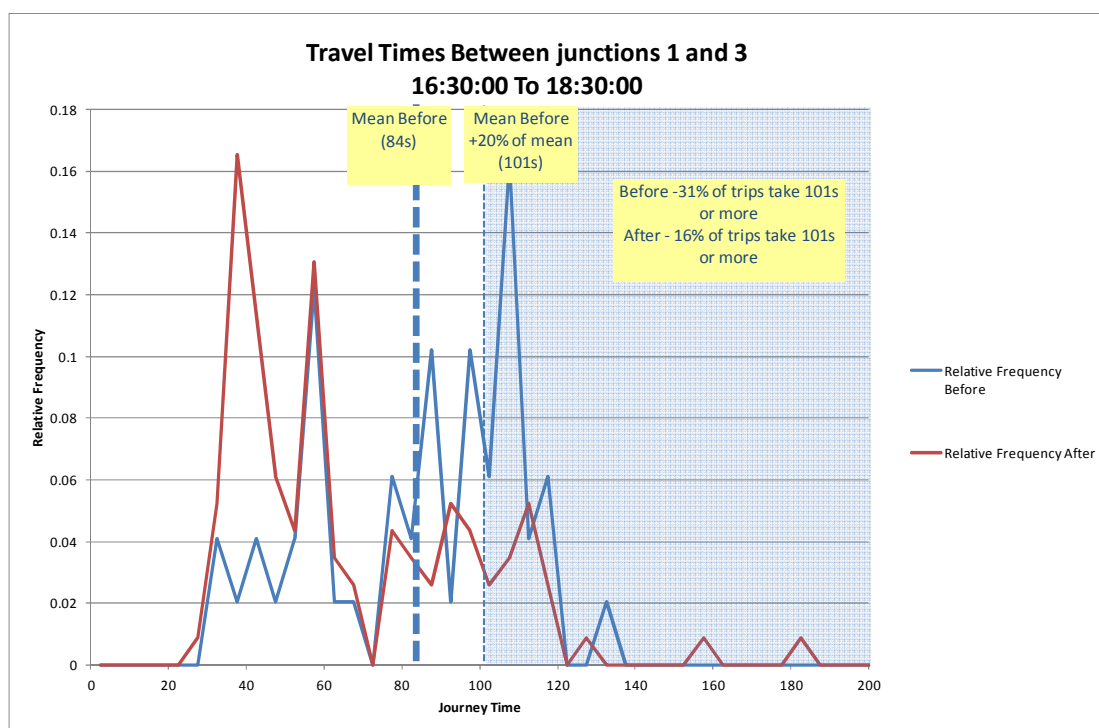
Scatter plot of travel times in the PM peak



This before and after data was further manipulated, to show an alternative representation in the form of a Relative Frequency graph (below). This similarly shows the 'before' travel time frequencies in *blue* and 'after' frequencies in *red*. This graph provides a better understanding of the relative bus travel times, and enables a test to be carried out which gives a measure of the proportion of buses which suffer 'higher' levels of delay (defined here as 'greater than Before Mean plus 20%'), the intention being to provide additional information about the longer bus travel time, not provided simply by the *mean* benefit.

The chart shows that the number of buses suffering 'higher' levels of delay (in excess of 101s in this case), fell from 31% in the 'before' situation to 16% in the 'after' situation.

Relative frequency graph of travel times in the PM peak



This Relative Frequency graph visually illustrates the changes made by the priority strategy. It shows a 'before' situation which has two distinct (blue) 'peaks' – a small one at around 60s and a much larger one at around 110s. This latter peak is an indication that most buses in the 'before' situation failed to get through the first green at Junction 3. In contrast, in the 'after' situation (red) the patterns are now changed – the peak at around 110s is much reduced, and there is a new distinct peak at around 40s, indicating that there are now far more buses being prioritized through the first green. This is a visual indication of a successful strategy outcome.

OVERALL BCR OF THE BUS PRIORITY SCHEME

At the commencement of the scheme the benefits of each junction were estimated, and an overall estimate of Benefit/Cost Ratio (BCR) derived for the scheme as a whole – with a value in excess of 8. Most of the junctions in the monitoring sample gave *actual* benefits greater than the estimated values, and these sampled junctions *alone* gave a BCR in excess of 2. Thus the key project metric has been achieved.

An assessment of the robustness of the scheme BCR shows that the final value is likely to be considerably higher than 2. In the unlikely event of the non-sampled sites ALL yielding under 75% of estimated benefits, there would still be an overall scheme BCR in excess of 7.

2) TRAMS - Strategic Priority in Edinburgh

The typically lower frequency of trams relative to buses, combined with greater passenger loading, makes the provision of high tram priority very desirable.

The Edinburgh Tram is due to go live in summer 2014 and will provide an efficient method of transport across the City Centre and out to the Airport. A key component of this is to ensure that the tram incurs minimum delay at traffic signals. This will be done through a mixture of STM and 'LRT' priority. STM will be implemented across the city centre network at 21 signal controlled junctions and crossings mainly controlled via UTC fixed time plans, whilst junctions further out towards the airport will be controlled with more traditional controller defined 'local LRT' priority.

Overall Network strategy


The key to achieving the aspiration of providing a high level of tram priority, balanced with minimal disruption to general traffic, is the ability to predict the arrival time of a tram at each junction. When coded with offline derived 'time of arrival' information (in terms of 'time in the signal cycle'), STM can then make strategic decisions on altering the downstream signal timings in advance of the tram's arrival at a particular junction. The object of this is to match the appearance of the green to the tram arrival, and in order to make it easier to achieve this at adjacent junctions it is envisaged that the network will be broken down into small coordinated 'sub networks', typically separated by tramstops.. In addition to this, STM will also monitor the greens, and calculate what compensation may be needed to minimize the disruption to other traffic.

Journey Time prediction

Since local LRT detection is being placed throughout the city centre (to allow 'LRT' priority to be used as a fallback mode to STM) STM will use this detection as its method of monitoring the trams.

The local LRT detector replies will be brought back via UTC reply bits, and can either be used in real time as logic inputs, or stored in the STM database for offline analysis of tram journey times. The STM priority strategies will use the detector replies to select the most relevant time-of-arrival values related to downstream junctions. The database will allow queries to be made between pairs of detectors by the Journey Time Tool.

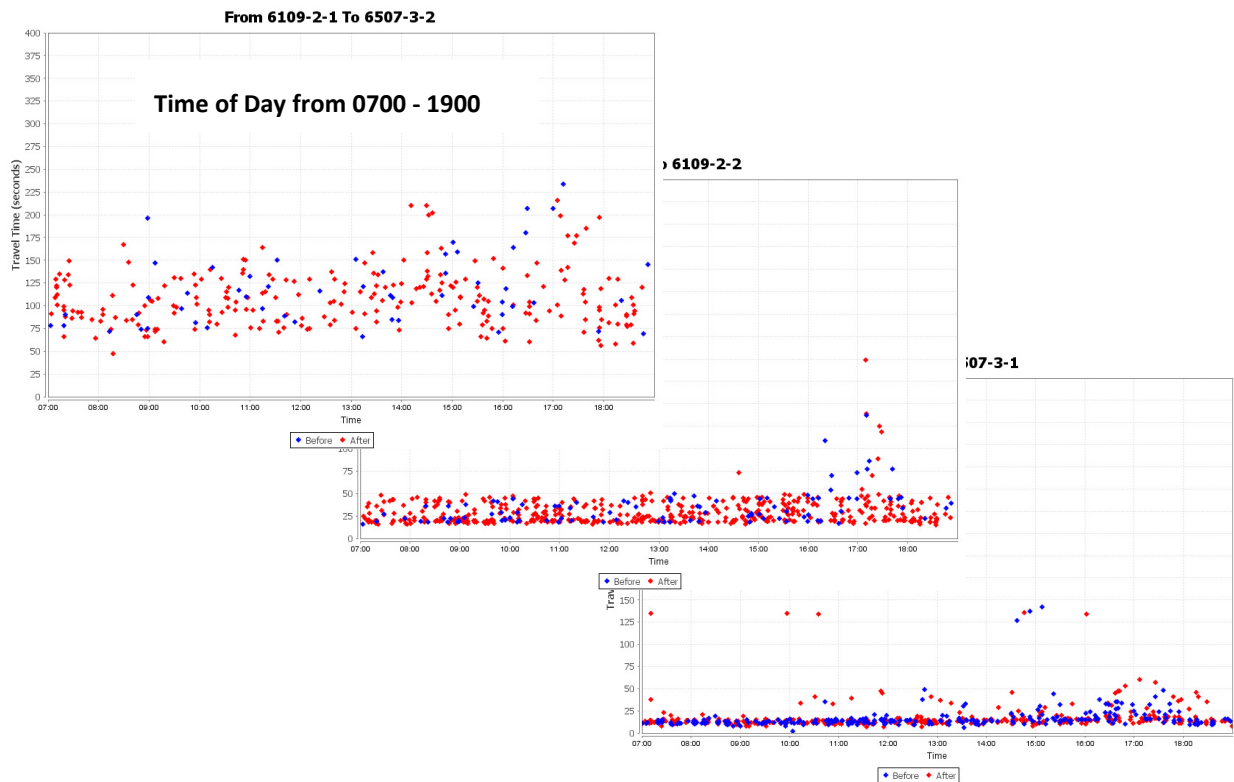
This diagram shows how detectors will have an associated journey time to downstream junctions



The Journey Time Tool will be used to assess where the journey time from upstream detectors is within tolerable limits of variability, such that a realistic decision can be made as to whether to shift downstream

timings 'early' or 'late'. This tool allows time distance diagrams and scatter plots to be produced, using either 'raw' data (as shown in the scatter plots below), or 'processed' data where any non-priority signal delay has been removed.

The Journey Time scatter plots below show the ('raw') Journey Time variability decreasing as each new detector is passed, and hence the distance to the target junction decreases



Tram Strategy Overview – 'Modes'

STM strategies use 'Modes' as a way of sub-dividing the various steps within a strategy. These modes have typically been used in bus priority to define basic extension/recall/compensation strategy steps. For Edinburgh Trams there will be 3 main priority 'modes' which will help define the strategy steps for an approaching tram (in practice these will be broken down into 'sub modes' to deal with such issues as '+ or -timing changes', and 'direction of travel').


Mode 0: No STM priority.

Mode 1: When a tram is detected on the approach to a junction(s), STM will access a set of pre-defined arrival times derived from the historic journey time database. When this estimated arrival time falls within reasonable levels of variability, small timing changes can be made at downstream junctions, and as subsequent detectors are passed the variability will decrease, allowing the timings to be refined – and larger changes to be made if required.

Mode 2: Once the tram is within the 'Prepare' detector on the final approach to the junction a similar style of priority to local LRT priority will be implemented.

Mode 3: On clearing a junction, STM will then handle the transition back to non priority **mode 0** and normal fixed time plans - ensuring no crash changes and applying any necessary compensation.

The diagram below shows a tram being detected a number of cycles before its arrival at the first of a 3 junction sub-network. Mode 1 is initiated and timings over the next few cycles are altered in anticipation of the trams time of arrival within the cycle. Mode 2 will then handle the final approach with Mode 3 running a compensation cycle



Mode 1 – the key mode


Mode 1 is the critical part of the priority. Mode 2 is necessary, but If Mode 1 has been successful then Mode 2 will only need to handle the appearance and clearing of the tram priority stage – resulting in a far less aggressive change in timings than with local LRT.

To enable Mode 1 to deal with both directions of tram, optimal signal timing for each sub-network will be pre-defined for each tram direction.

Whilst the ‘upstream’ junction in a sub-network is being altered to match the time of arrival of the tram, the downstream junctions within the sub-network will be moved to the pre-defined optimized offsets, and thereafter moved together as a group until the tram passes. The key is to balance the need for quick timing changes with the likely disruption to general traffic.

A necessary part of Mode 1 will be directed towards the issue of whether timings within the sub network are required to be made ‘early’ or ‘late’.

The diagram below shows pre-defined offsets within a sub-network for Eastbound and Westbound trams.



3) Fire – Green Wave Pilot

During 2010 the West Yorkshire Fire & Rescue Service (WYFRS) became aware of the West Yorkshire ‘Traffic Light Priority’ project, and approached Metro with an idea of adapting the system for use on their Fire Appliances.

Although by law Fire Appliances can proceed through red signals with caution, it can cause difficulties for both the firefighters and other motorists, and naturally slows appliances down as they attempt to safely negotiate other traffic. In addition to this ‘in junction’ delay, sometimes appliances cannot get to the signals quickly because of queuing traffic on the approaches to junctions.

Appliance accidents, and response times

In 2009 the Fire Service undertook some research into accidents and ‘999’ response times. The research showed that in 2009/10 there were 2922 accidents involving Fire Appliances – with around 40% of these being whilst they were ‘on blue lights’ (data compiled and produced by the Chartered Institute of Public Finance and Accountancy (CIPFA).

This research also highlighted an increase in response times in line with increasing growth in traffic levels:

West Yorkshire Data, average dwelling fire response time(min) and average traffic levels million vehicle/KM

	1996	2006
Response time	5.4	6.3
Traffic levels	13,868	15,818

Using response time fatality rate relationships WYFRS predicted that the increased response time may contribute to around:

- 13 additional fatalities in dwelling and other building fires each year
- Possibly 65 additional deaths in RTAs
- An £85 million increase in building fire damage.

LEEDS GREEN WAVE PILOT (DEC 10-MAR 11)

WYFRS was driven by the desire to increase road safety for their appliances and other road users, and also to reduce appliance response times. In consultation with Leeds City Council, they established a set of criteria for the success of the pilot.

Pilot success criteria

- Give a green light to each emergency journey in at least 90% of instances,
- Ensure and increase safety levels for the public whilst using the road during an emergency journey,
- Flush traffic from both the next and subsequent junctions ahead of the Fire Appliance's arrival.

Hunslet Fire Depot in South Leeds was chosen for the trial. This Depot has 3 main routes that Fire Appliances could use to access different parts of the city, all suffering from congestion in the peak periods and having a number of signal controlled junctions. Two Fire Appliances were fitted with the same AVL priority hardware as used by buses, the units becoming active once the Appliance is under 'blue lights'. As the technology was already in use, adapting it for Fire Appliances was relatively straightforward.

STRATEGIES

Because Fire Appliances are low-frequency vehicles, it is possible to give them 'high priority' - because lives are at stake it doesn't matter if other traffic is inconvenienced by short-term queuing.

It was agreed that Appliances in Leeds could be granted the highest level of priority possible, with certain guidelines being set out to help define this:


- Stage skipping could be considered if deemed safe to do so.
- If stage skipping occurred to the detriment of a pedestrian stage, the demand must be satisfied on the next stage move following priority.
- Compensation was not necessary due to the infrequent disruption, and the understanding by the public having just seen a fire appliance 'on blues' proceed through the junction.

As well as the low frequency of Fire Appliances compared to buses there are other key differences that needed to be considered, which the strategies had to cater for:

- **Advantage:** Fire Appliances do not stop at bus/tram stops. Analysis of the travel times along the routes showed more consistency than buses, although there were still time-of-day trends in line with peak hour congestion.
- **Disadvantage:** Fire Appliances do not follow known routes. On leaving the depot all routes had to be activated with priority until it could be determined which route the Fire Appliance was taking. Appliances can also turn off the route unexpectedly, or switch off their blue lights.


Hunslet Southern Route

Below we see one of the three routes from the Hunslet Depot. The Fire Depot itself exits into a signalized junction (Node 1) which uses a traditional push button and hurry call stage at the signals. This push button is the first 'trigger' in initiating the priority, but at this point the route is still unknown and so the first junctions on ALL three routes are activated (including Node 2 in the route shown below). Only once the Appliance commits to a particular route can the two unused routes be cancelled and the full strategy on the selected route be committed to.



Southern Route – 4 junctions and 1 pedestrian crossing

Numerous AVL Triggers are used to constantly monitor the Appliance's progress and to determine its route



EXAMPLE OUTLINE LOGIC

As a Fire Appliance passes along the route its progress is tracked by AVL triggers placed around every 50m. Having numerous triggers gives some redundancy, especially at the exit from the Fire Depot where there were concerns over how quickly the Fire Appliance's onboard GPS would be able to start tracking its location. The triggers are often grouped together to initiate the same strategy, for example southbound triggers 1-5 all initiate the same strategy at the downstream junctions. Not until triggers 6-8 are reached does the next strategy get initiated, and so on along the route. Below is some example high level logic for 2 example triggers;

Trigger 1 – at the Fire Station Exit

- Node 1 – Travel time =0s
 - Hurry Call stage calls all red at Fire Station junction
- Node 2 - Travel time ≈20s
 - If in Priority stage then hold stage
 - If not in priority stage then immediate move to priority stage (allows stage skipping)
- Node 3 – Travel time ≈45s
 - If in priority stage 1 then hold stage.
 - If in stage 2 then immediate move to stage 3
 - If in stage 3 at start of green run min and move to priority stage 1
 - If in stage 3 and already run stage min then do nothing
- Node 4 – Travel time ≈60s
 - Do nothing
- Node 5 – Travel time ≈120s
 - Initiate a plan to clear queuing traffic (more plan aggressive in PM peak)

Trigger 11 – at the exit of Node 2

- Node 2 - Travel time 0s

- Smoothly drop back to UTC fixed time plans ensuring pedestrian stage run next.
- Node 3 – Travel time ≈15s
 - If in priority stage 1 then hold stage.
 - If in stage 2 then immediate move priority stage 1
 - If in stage 3 then immediate move priority stage 1
- Node 4 – Travel time ≈30s
 - If in priority stage 1 then hold stage.
 - If in stage 2 then immediate move to priority stage 1
- Node 5 – Travel time ≈90s
 - Run plan to clear queuing traffic.

It can be seen that as the Fire Appliance moves along the route the urgency in how a junction moves to the priority stage is increased. By initiating those junctions with sufficient pre-emption early enough, it is possible to avoid harsh stage changes and stage skipping.

Other supporting logic

- The clear-down of the priority stage must allow sufficient time for a following Appliance.
- Any trigger fired on a side road will cancel the green wave – this indicates the Appliance leaving the route
- A timeout of 30s from the last trigger received will also cancel the greenwave - this covers the Appliance turning off its blue lights.

PILOT RESULTS

The overall reliability of the green wave pilot was between 85%-95%, with the higher values being achieved towards the end of the pilot.

Unlike buses and trams it is difficult to fully quantify the benefit to Fire Appliances in terms of pure delay reduction, since they routinely 'jump reds'. However, evidence from WYFRS firefighter testimony indicates that on the approach to an active junction there is significantly reduced or zero standing traffic, which is supported by the following quantified delay benefits:

Three Routes Activated (delay in seconds)

Route	Average			Worst Case		
	Before	After	Saving	Before	After	Saving
Southern Route	90	77	13 (14%)	240	89	151 (63%)
Northern Route	59	54	5 (8%)	90	66	24 (26%)
Garnet Rd Route	35	30	5 (14%)	41	33	8 (19%)

In addition to the delay benefits ranging from 19-63% in the pilot, there is also an anticipated reduction in accidents costs, which will require to be monitored over a longer time period.

WYFRS consider that the Green Wave system will mean safer and faster journeys for firefighters en route to emergencies. It should also assist other drivers at traffic signals, who are known to be unpredictable when faced with an emergency vehicle using blue lights, by allowing them to continue rather than queue. This view is reinforced by the fact that in 2011 the WYFRS Management Board approved an expansion of the pilot to include two more depots and 4 more Fire Appliances in Leeds, with a view to it being rolled out across West Yorkshire.

OVERALL CONCLUSIONS

Three strategic priority examples have been discussed in this paper, with increasing priority needs, but all making use of the STM priority system. Two of these are implemented on the ground and one (tram priority example) is under active development:

- **Bus Priority** – the implementation of complementary priority elements at different junctions along a route. This has demonstrated that the targeting of particular delay issues with specific techniques can yield significant benefits.
- **Tram Priority** – the development of a more gradual system designed to manipulate junction timings over a longer route. It is anticipated that the use of longer time-horizons will enable the achievement of minimal tram delay without excessive traffic disruption.
- **Fire Appliance Priority** – the piloting of an ‘intelligent’ green wave which clears out traffic in advance of arrival of a fire appliance and reverts to normal operation at the earliest opportunity. The priority pilot is seen as being very successful by the Fire authority, who intends to roll out the system more widely.

References

Reference 1: SPRUCING Up Your Bus Priority; Brent Collier, Sheffield City Council; JCT Symposium; 2009.

Reference 2: Priority and programmability - with STM-SPRUCE; Mervyn Hallworth, Leeds City Council; JCT Symposium; 2006.

Reference 3: High-performance bus/tram signal priority; Mervyn Hallworth, Leeds City Council; JCT Symposium; 2004.